2010

What Makes a Belief Warranted? A Pragmatist's Answer

Luke Herrine
Oberlin College

What Makes a Belief Warranted?
A Pragmatist's Answer

Luke Herrine

INTRODUCTION

It is natural to think that some things are better to believe than others. It is equally natural to think that we should believe the good beliefs and not the bad. In our everyday lives, we rarely have any trouble distinguishing the former from the latter. One important question to ask, though, is what makes good beliefs good and bad beliefs bad. Epistemologists have long tried to find a principled distinction between good and bad beliefs. I think, though, that most epistemologists have gone about finding such a distinctions in the wrong way.

Attempts to answer the question of what makes good beliefs good have generally come out in theories of *warrant* and *justification*. These terms are often used interchangeably and their definition can vary from author to author, so it will help to say how I shall deploy them. I will use 'warranted' to mean epistemically good—that is, something we should believe—and I will use 'justified' to mean supported by an appropriate bit of reasoning. Thus, all justified beliefs are warranted but it is at least arguable that not all warranted beliefs are justified.

We can accordingly rephrase the question of what makes good beliefs good as what makes a belief warranted. Some epistemologists may have proposed their theories of warrant without the former question in mind, however. They may have seen themselves as performing a conceptual analysis of the concept of warrant or of looking for the missing piece to Gettier's puzzle about knowledge and true belief. Though I think that either of these goals should be seen as merely instrumentally useful in service of discovering what makes something good to believe, I do not want to pick that fight here. I shall simply take from theories of warrant what is relevant to creating a principled distinction between good and bad beliefs.

It has long been thought that the prime example of something we should believe is something we know. Since Plato, justification has been commonly thought to be the essential

quality that a belief must have in addition to truth in order to be considered knowledge.[i] As such, for much of epistemology's history, reflections on the nature of justification have focused on how it underscores knowledge attributions. After Gettier's famous essay[ii] challenging established thinking on justification's role in knowledge attributions, the focus on the relationship between knowledge and justification became all the more intense. The notion of warrant as something slightly different than justification arose in this flurry of responses to Gettier's challenges. Many epistemologists continue to analyze reasons for accepting a belief in terms of that belief's relationship to knowledge. But, as exemplified in work by Timothy Williamson,[iii] dividing knowledge up into constituent parts has fallen out of favor with many others.

Even after setting such conceptual analysis aside, it remains important to try to understand what it takes for a belief to be warranted. On the one hand, Williamson argued for a cessation to the breaking of knowledge into component parts because he wanted to focus our attention on knowledge as a primitive concept. His position affirms even more strongly a reliance on knowledge in figuring out what we should believe. [1] My position will follow his lead in cleaving warrant from knowledge, but I will do so to attempt to find a more fruitful explanation of the former rather than the latter. I think it is much more intuitive and certainly more practically useful to ask questions about the nature of what we should believe than to ask questions about the nature of knowledge. My reasons for this position should become more evident as we move forward, so I do not wish to belabor it now.

Another historically relevant bugaboo attached to the concepts of warrant and justification is certainty. Mathematically inclined thinkers from Descartes to Russell have suggested that our only warranted beliefs are those of which we are certain. Contemporary

---

[1] In Williamson's "knowledge-first approach", we must know what to call knowledge in order to know which beliefs have warrant. One's beliefs gain their warrant in virtue of evidential relations with one's knowledge. I think his view is unsatisfactory for some reasons that will come out below.

philosophers find little attraction to this view, so I will simply treat it as defunct. Rather than asking of what can we be certain, the more tractable question seems to be: what quality of a belief makes it fit our standards? Phrased another way: what makes some things reasonable/rational to believe and some things not? Many attempts have been made to answer this question. I will explore internalism in its foundationalist and coherentist guises, externalism in the forms of reliabilism and Plantingan proper functionalism, and the evidentialism explicated by Conee and Feldman (with a nod to Williamson).[2]

Though all of these views have much to praise, my focus will be on the problems that each faces. The sorts of problems I will point to will dig up an underlying issue common to all of these approaches. Briefly stated, epistemologists have generally ignored the issue of warrant attribution. Though epistemologists have had many interesting things to say concerning what about a *belief* or a *believer* affects warrant they have had little to say concerning what about the *attributor of warrant* affects warrant. They have missed epistemological parallax, or so I will argue.[3]

Confronting the problem of epistemological parallax will be the goal of my positive account. Facing up to the possibility that facts about *us* factor into which beliefs we count as warranted opens a can of worms. We will no longer be able to take for granted our own epistemic intuitions as obvious facts about what makes some beliefs better than others. We should not rely on any guarantee that everybody will share our intuitions or that our intuitions are invariant. I will explore in some depth what goes into the process of warrant attribution while

---

[2] Conee and Feldman defend an internalist version of evidentialism and Williamson an externalist version. I treat Conee and Feldman's views apart from other internalist views because of important differences that I will address later.

[3] The recent popularity of contextualism about knowledge has made epistemologists more aware of these factors. I will not address contextualism head on here, mostly because it deals with *knowledge attributions* and not *warrant attributions*. It may be worth noting, though, that my position is much more closely aligned with the contextualism of Michael Williams than that of Keith DeRose.

taking into account both attributor and believer factors without taking for granted attributor's intuitions or relying on their invariance. My exploration will result in a sort of pragmatism: pragmatism about warrant.

Various thinkers have anticipated several of the explanations I shall propose. I borrow liberally. My purpose is most often to borrow concepts and glean inspiration rather than to critically engage with a work or an oeuvre, so I will ask readers to excuse what may occasionally seem idiosyncratic readings. Additionally, a few of the thinkers I reference will be unfamiliar to or unpopular with contemporary analytic philosophers. I do my best to extract only the relevant parts of these works and to make them comprehensible. Unfortunately, my views may seem strange due to this borrowing, but I hope I can do the appropriate work to make them at least defensible. In this vein, I follow up my positive account with a discussion of how I think it solves the problems that I posed for other views and with responses to objections. I will conclude with a discussion of the implications of acceptance of my proposal. Like some pragmatists before me, I will suggest that epistemology should look decidedly different. But first, I want to explore how it looks now.

TRADITIONAL INTERNALISM

Internalism has been the dominant (often thought the only) epistemological position for most of the history of epistemology since Descartes. A distinguishing feature of this view is the identification of warrant and justification: the internalist focuses on how beliefs fit together within a given subject's doxastic architecture. In order to understand why a given belief has warrant, the internalist would say, one must see how it fits with other beliefs held by the believer. One must understand what reasoning the believer possesses for a questionable belief in order to evaluate it. Internalism often goes hand-in-hand with what Michael Williams has called epistemological realism. He uses this term to indicate the position that there exists some single mostly-invariant epistemic structure to our beliefs, which we uncover when we engage in their justification. So a belief gains its warrant from underlying justificatory connections with other beliefs. When engaging in the defense of a given belief of ours, we must merely unveil our pre-existent justificatory structure to demonstrate its warrant. Because the structure itself provides the belief in question with its warrant, warrant and justification prove indistinguishable.[4, iv]

By far the most widely accepted (both historically and contemporaneously) internalist position is foundationalism, which posits a certain set of basic beliefs that act as ultimate justifiers to any non-basic belief within a person's belief system. So the warrant of any given belief derives from its relationship to beliefs more basic to it on down until the basic beliefs. Basic beliefs enjoy the privileged status of needing no other belief to justify them. Various possible reasons have been proposed for their special status, but I will not address them here.[v]

---

[4] Basic beliefs might be warranted without being justified according to the internalist foundationalist. How this works without some sort of externalism would be difficult to say. Foundationalists in the first half of the 20th century struggled mightily with this problem. For a review (and a critique) see Williams's *Groundless Belief*.

Most modern-day internalist foundationalists (for brevity, I shall omit "internalist" for the remainder of this section) hold the empiricist position that basic beliefs gain their warrant from the fact that they derive from faculties that we justifiably trust. In fact, many contemporary foundationalists use a bit of externalism to get at the warrant of basic beliefs, but for now we shall leave this notion unquestioned.[vi] It will not be important to get into these details for the sort of critique in which I want to engage.

Foundationalism lends itself especially well to structural metaphors for belief systems. The position's name itself leads the way: basic beliefs serve as the *foundation* upon which one builds the rest of one's doxastic structure.[5] The stability of any given belief depends on the other beliefs that support it, which ultimately depend on the stability of the foundation of basic beliefs. Though we have beliefs about a wide range of topics that may not be directly connected to each other, all of them are ultimately built upon this foundation. If the warrant of some belief comes into question, all we must do is demonstrate the stability with which we have built the doxastic structure that supports it. If challenges persist, we can always eventually find ourselves at our foundational beliefs, which rest on undeniably solid ground.

The other commonly accepted internalist position is (the internalist version of) coherentism. Coherentists might be described as having absorbed the Quinean insight that no belief faces the tribunal of experience singularly but only as part of an entire belief system.[vii] As such, the warrant of a given belief derives not from its relationship to some subset of basic beliefs that have a special relationship with experience or the external world, bur rather from its relationship to a believer's whole doxastic structure. As long as the justificatory relations that bond beliefs together are good, any given belief is good. Though no agreed-upon complete list

---

[5] Sosa's metaphor of a pyramid in "The Raft and the Pyramid" seems to me misleading, because it makes the foundation the broadest part and it has only one apex. Not that these issues lead to any particular problems with his analysis.

exists, belief-bonding material is commonly thought to be made of such things as logical entailment, Bayesian probability relations, explanatory relations, etc.[viii] There is no such thing as modifying just one belief, because every belief is bonded to a whole set of other beliefs (and, depending how thorough the holism, to *every* other belief) through such relations. Thus, whenever it seems that just one belief is under review, actually a whole group of beliefs and their relationships face rejection. The crucial factor in determining warrant/justification is how beliefs cohere with each other rather than on what they are based.

The coherentist picture of beliefs resembles a web (or a raft, according to Sosa). Beliefs are held together by justificatory threads. Knock off one belief and a whole series of others fall with it. A section of beliefs can be repaired, but never just one. Justification is non-linear—beliefs can be mutually supporting without being viciously circular. Just as a spiderweb depends for its support on outside structures on which it can stick (the crook of a tree or the corner of a doorjamb, for instance), a belief system depends on the outside influence of experience for its support. However, not just the beliefs at the outer edge of the web are affected by experience. "A conflict at the periphery occasions readjustments at the interior" of the web, Quine said[ix] (though I quote him in support of this position, Quine should probably not be considerd an internalist coherentist). The problem may not be with the peripheral beliefs, but with their relationships to interior beliefs.

Both of these internalist positions have plenty of *prima facie* appeal, especially, I suppose, because of their usage of metaphor. Visual animals and metaphorical thinkers that we are, such theories draw us in.[x] But why they are intuitively appealing is incidental to the present discussion—a topic for experimental philosophy, perhaps. In any case, they certainly have other redeeming features. More incumbent on our pursuit of the best theory of warrant is why we

should not adopt either internalist position. Of course the foundationalist and the coherentist each has her reasons for rejecting the other's position, but these will not concern us. The bigger problem lies with the more encompassing commitment to internalism.

Recall that we are viewing internalism as an analysis of the principled difference between good and bad beliefs. To do this explanatory work, both the coherentist and the foundationalist appealed to global features of individuals' doxastic structures. But why should we assume that all beliefs gain warrant by the same means? Certainly that would make things easy for us theoreticians—we would just have to figure out the rules for good beliefs and then apply them whenever a specific case of questionable belief came up. Such a position makes especial sense if one thinks about it in the context of the enlightenment in which it was developed: reason was supposed to be able to solve everything, the unitary self (soul) was widely accepted, and metaphysical systems were common ways of explaining each individual fact or belief. But nowadays things do not seem quite as simple or easy as our philosophical predecessors may have hoped. Scientists' way of approaching things certainly seems quite different from that of historians or art critics or monks, and even within science, molecular biologists do not go about obtaining beliefs in quite the same way as physicists or economists. We go about things differently depending on what we want to explain or whether explaining is our goal at all. Even if all of these methodologies were merely different decorations for a more basic unitary structure of beliefs, we would need to understand why it looks so different in these distinct dresses. Moreover, that such a unitary standard exists should be a result and not a starting point in explaining how warrant works. "All we know for sure is that we have various practices of assessment [of warrant], perhaps sharing certain formal features; it doesn't follow from this that the various items given a positive rating add up to anything like a natural kind."[xi] Much more

work would need to be done to find or argue for such a unitary standard, and this work would need to be done *before* internalist answers to why we should believe some things rather than others would have any explanatory force.

As already mentioned, internalism also conflates justification with warrant. When asked "what is it for a belief to be warranted?" the internalist hears "how do we justify a belief?" and in the same action assumes that what comes to light in justifying behavior is what underlies a belief's being warranted in the first place. Not that this is an unreasonable first move. When we ask ourselves the question of what we should believe, of what makes a belief warranted, an obvious first place to look for answers is how we defend our beliefs. The problem arises when we take these defenses to be exercises in introspection.[6] Belief defense depends as much on one's concept of one's audience as one's concept of one's own beliefs. When we engage in justification, we aim to absolve ourselves of epistemic blame. Doing so might have little to do with reflecting on our own belief-gaining mechanisms. One would be ineffective in defending her belief if she did not take into account the standards of the challenger.

The confusion between warrant and justification may also have arisen because if I say "belief S is justified" outside of any specific context then it is unclear whether I mean "belief S was questioned by X and appropriately defended by Y" or "X's belief S is warranted ". In other words, outside of any context it is unclear whether the participial form "justified" is the passive perfect form of the verb "justify" or merely functioning as an adjective—whether it means "well defended" or "warranted". But no matter why the conflation arises, it is important to note that it *is* a conflation. Given the above considerations, we seem to have little reason to think that these

---

[6] And making such a move moves one to Bonjour's absurd invocation of the "doxastic presumption" in his defenses of coherentism (op. cit.). His own (later) reasons to reject this principle are perhaps the most damning. They can be found in Laurence BonJour, "The Dialectic of Foundationalism and Coherentism" in *The Blackwell Guide to Epistemology* ed. John Greco and Ernest Sosa (Massachusetts: Blackwell, 1998): 126-130

two possible parsings of the term "justified" reflect the same underlying phenomenon. The great majority of the time, our beliefs are taken as warranted without having to be justified (i.e. defended), and how they are justified may have little to do with their warrant.

Internalism's most glaring and most discussed problems is its uncomfortable relationship with skepticism. The very invocation of a foundation or a global coherence in the first place can be instructively conceptualized as a response to Agrippan skepticism. To briefly summarize, the Agrippan skeptic worries that once we confront a particularly pesky challenger (imagine a 3-year-old, for instance, repeatedly asking "why?"), we face one of three seemingly unsatisfactory options, known as Agrippa's Trilemma. First, we could be unable to find any stopping point for justification whatsoever and continue *ad infinitum nauseamve* with no satisfactory response for the challenger. Second, we might use some circular reasoning; i.e. we might use a belief for justification that has already been challenged. Third, we might arbitrarily choose one belief at which to cease justifying. The foundationalist might be thought to have discovered a fourth, more palatable option: we end justification on a basic belief. Because these basic beliefs have privileged status, it would not be arbitrary to end on them--they just are the foundation upon which all of our beliefs rest.

But what if these foundational beliefs are not as immune to challenge as the foundationalist thinks? The Cartesian skeptic, for instance, has been thought to call into question the certainty (or, in a more powerful modification, the legitimacy) of foundational beliefs. If the foundationalist cannot fend off these attacks, then she finds herself in dire straits indeed, for she has built her epistemological edifice so that if the foundation is unstable, then the whole structure falls. Because the foundationalist must presuppose that all beliefs depend for their justification on basic beliefs, she must maintain that unjustified basic beliefs would render every belief

unjustified. It is in the foundationalist's best interest to fend off the Cartesian skeptic, because if the skeptic is not refuted head-on, then the foundationalist must concede that none of our beliefs are justified by her standards.

Unfortunately for foundationalism, nobody has been able to find a good direct counterargument to the skeptic. Even Kant, who called it "a scandal of philosophy" that the existence of the outside world had not yet been proven by philosophers,[xii] failed to provide a generally accepted knock-down counterattack to the skeptic. Indeed, much of the history of modern epistemology has resembled Reagan-era discretionary spending in that an enormous amount of resources have been devoted to defending our beliefs from the impending doom threatened by the apparent arsenal of the looming skeptical enemy. Yet still no foolproof defense has been worked out.[xiii] The problem is that if we think that internalist foundationalism explains what makes our beliefs justified, we would need a foolproof defense from the skeptic to be able to have *any* justified beliefs.

The coherentist provides a different response to the Agrippan skeptic. Rather than focusing on a special class of beliefs, we should end the regress by understanding why circular reasoning can be virtuous. We should jettison questions of regress in favor of seeing the trees for the forest. But this response confronts another sort of skepticism, which can be stated quite simply. Even if all of our beliefs cohere perfectly well, they could have nothing to do with the real/external world. Because we want our beliefs to correspond to, or at least be accurate descriptions of, the external world, coherence does not seem to be enough of a criterion for warrant. A totally deluded person can have perfectly coherent beliefs without having any accurate beliefs.[7]

---

[7] See also Plantinga's case of the hiker with a frozen visual field in *Warrant and Proper Function*, 179.

Both internalist positions prove especially weak against skepticism because of two basic features of the versions of internalism explored here. The first is the appeal to global features of an individual's belief system. Skepticism exploits the reliance on a single set of factors for the warrant of all beliefs by calling into question the legitimacy of those very factors. If the warrant of any given belief depends on very general features of one's doxastic structure, then challenges to the general features call all specific beliefs into question. For the foundationalist, questioning the legitimacy of basic beliefs leads to potential problems with all beliefs. For the coherentist, pointing out that perfectly coherent belief systems can be wildly inaccurate potentially problematizes all of our beliefs. Of course it is also damning that these global features describe the *internal* structure of beliefs, because both Cartesian skepticism and anti-coherentist skepticism challenge the metaphysical relationship between the internal and the external world, where the former is supposed to represent the latter. If we have no mechanism of knowing whether the internal is in touch with the external, then it remains an open possibility that even the beliefs best justified by internalist lights in fact have no relation to external reality.

EXTERNALISM

Dissatisfied with the inadequacies of internalism, various authors have developed theories of warrant that they call, contrariwise, "externalism". To answer the question of what makes something good to believe, externalists move the focus away from how an agent's beliefs fit together (though they may have some independent interest in those issues) and use various objective facts about the believer in relation to an environment to determine the warrant of a belief. They propose that "warranted" be seen as an evaluative term in epistemological discourse, like "good" in moral discourse. In each case we evaluate a human action according to the standards of a given discourse or set of discourses. While internalists had asked *how* something is justified, an externalist thinks that the focus should turn to *why* something is, in actual fact, warranted. In order to do so, we are directed to look at beliefs that we think are well formed and then figure out what they have in common that renders them thus.

Reliabilism is the most enduring externalist theory. It hones in on *the manner in which* beliefs were formed. Reliabilists hold that if a belief arises from a reliable belief-forming process, then it is warranted. For a belief-forming process to be reliable, it must be the sort that mostly produces true beliefs with the given inputs. Which belief-forming processes are reliable are said to be vaguely delimited to reflect our messy everyday attributions of warrant. Some belief-forming processes produce true beliefs nearly always while some just do so more often than not. The former should certainly be considered reliable while the status of the latter may be up for debate. To a similar end, reliability is posited to be a matter of degree—some belief-forming processes are more reliable than others without any bright line separating which engenders warrant and which does not. Thus, seeing a pie in the window from a foot away with

normal 20/20 vision is more reliable than seeing a pie in the window from 30 feet away, which is more reliable than seeing a pie in the window from 200 feet away. At some (vaguely established) distance, 20/20 vision is no longer reliable, and, as such, a belief produced from this faculty no longer has any warrant.[xiv]

Reliabilism maintains a foundationalist flavor, as it designates a certain set of immediately warranted beliefs that provide the basis for the warrant of the rest. Immediately warranted beliefs arise from belief-independent processes (i.e. they are not generated from other beliefs—perception and memory work this way) while mediately warranted beliefs arise from belief-dependent processes (i.e. they are based on other beliefs). A belief-dependent belief can only be called warranted if the belief-independent beliefs on which it depends and the cognitive operations that led the subject from these beliefs to the belief-dependent belief are reliable. However, this warrant-bestowing process does not depend solely on the present state of the subject's doxastic architecture, as the internalist foundationalist would have it, but rather on the diachronic ontogenesis of any given belief. The process is what is important, not the believer's state of mind.

Alvin Plantinga has offered another version of externalism. On his account, "a belief has warrant if it is produced by cognitive faculties functioning properly in a cognitive environment congenial for these faculties, according to a design plan successfully aimed at truth".[xv] If any aspect of this formulation fails to apply, then the belief does not count as warranted. The details of this formulation could take up a book—and have done so—so my account will be necessarily abbreviated, but the general idea is as follows. Some of our cognitive faculties have been designed (by God, evolution, or whatever) to provide us true beliefs in a given set of

circumstances. These beliefs count as functioning properly when they behave according to the truth-directed design plan.

To get at what Plantinga has in mind here, it helps to compare a human to a car. A car is designed to transport cargo over certain types of surface by means of a combustion engine. If the alternator fails and the car cannot run, then it is not functioning properly—it does not accomplish the goal for which it was designed. If the car fails to transport cargo across a mile of quicksand, then it should not be said to be functioning improperly, but rather functioning properly in an environment not congenial to its design. If an engineer designs a car to make its cargo explode rather than transport it, the car is functioning properly when it does explode the cargo, but it functions according to a design plan not aimed at transporting cargo, the proper aim of a car. Similarly, some of our faculties have been designed to produce true beliefs[8], and these beliefs count as warranted if they function the way they are supposed to in the sort of environment in which they are supposed to function well.

Externalism—in the two leading forms discussed here, but also more generally—has a general procedure for uncovering what make something good to believe. First, the externalist finds methods of belief acquisition that she trusts/epistemically values. Putting it this way is a bit unfair to the externalist, because the epistemically valuable beliefs are not supposed to be idiosyncratic to the particular externalist in question, but rather indicative of what any reasonable person would value. The way such beliefs are usually delimited is by tacitly appealing to the commonsense of the reader: "look, here is an obvious case of a warranted belief". The possibility that one could think otherwise does not even come into play. In any case, the general idea is to find types of belief that have epistemic value (though *for whom* goes unmentioned). After

---

[8] Some cognitive faculties might be said to have been designed for two sorts of things, exhibiting *functional multiplicity*, but it is only the part of the cognitive faculty's design that points towards truth that matters to epistemologists (ibid, 25).

finding these beliefs, the externalist analyzes what about their methods of acquisition we value

with an eye for some shared valuable property. Generally this has been done by showing that the

valued methods lead to true beliefs in the appropriate conditions and then detailing the specifics

of the methods' means of producing true beliefs.[9]

It is the first step with which I want to take issue here. I find the thought unavoidable that

in pointing out which beliefs we praise and which we criticize, all we are discovering is what our

epistemic preferences are. Not that this information has no value. But its value is not the

discovery of the best ways of going about things, it has at best discovered what our

commonsense intuitions tell us are the best ways of going about things. These facts have

importance insofar as they tell us about what sorts of epistemic judgments people who think like

us[10] make. However it would be overly conservative and ethnocentric to think that *we* do the best

job making such judgments and thus that the nature of warrant can be discovered only in our

own epistemic evaluations. [11]

One way to see the problem here clearly is by confronting the existence of cognitive

diversity. Stephen Stich expresses this confrontation especially clearly and urgently.[xvi] He notes

that the fact that we *could* reason quite differently paired with the fact that many other people

actually do so leads us to question whether we have any particularly good reason for doing it our

way. Because (at least many of) our cognitive processes come to us as part of our cultural

inheritance and others have different cultural inheritances, we should question whether the

cognitive processes that happened to be passed down to us are the best for doing the job we want

---

[9] This procedure also goes for other forms of externalism not discussed here (Nozick's and Armstrong's, for instance).

[10] And who is "us", here? I generally mean highly-educated western scientifically-minded 20th- and 21st-century philosophers. In other words, those who have understood and been sympathetic to the works of externalist epistemologists.

[11] Goldman acknowledges the tacit ethnocentrism to some extent on p. 104 but he does not examine its implications.

them to do. The externalist response might be that these different cognitive processes—however they were inherited—can all be analyzed according to the same standards of warrant. We can find out if they are the best for doing their job by testing them against standards of warrant. But the standards of warrant that the externalist employs are merely those of her own culture's commonsense, and "surely the evaluative epistemic concepts embedded in everyday thought and language are every bit as likely as the cognitive processes they evaluate to be culturally acquired and to vary from culture to culture".[xvii] In light of the evaluative diversity that surrounds us, adopting standards that strike one as commonsensical seems a somewhat arbitrary move.

A possible externalist rebuttal might go something like this: "All I am saying is that we want to adopt methods of belief acquisition that lead us to true beliefs. Even if another culture goes about things differently, we can analyze its methods of belief acquisition in terms of how well they produce true beliefs. Regardless of which culture a given epistemic practice comes from, it can be evaluated in terms of its ability to lead its practitioners to true beliefs." Unfortunately, such a response only pushes the problem back a level. Different people have different concepts of which beliefs are true. These differences can turn on quite deep disagreements (such as the question of whether God exists or whether the earth is round) or those more trivial (such as whether the concept of water includes XYZ in addition to H_2O). Setting an absolute standard based on only the set of beliefs that one currently thinks is true seems to turn on the same arbitrariness as before.

To emphasize this point I think it is helpful to think about what people thought was true over various periods of history rather than across present cultural boundaries. It is especially helpful to think about the history of scientific discovery. In any scientific field what scientists think is true changes relatively often and can at any given time be up for debate. Additionally,

many sciences periodically change standards for what count as good lines of reasoning. If at any point in the history of science, a scientist were to stop and say: "okay, now we have not only the truth of the matter, but also the best way of evaluating whether or not we have gotten to the truth of the matter", we would say that that scientist was being overly epistemically conservative. To see this, imagine if Carl Linnaeus had said such a thing, or Charles Darwin. Evaluating our own relation to truth so positively in the present day would also seem to render us guilty of myopic epistemic conservatism we would attribute to them. We should leave open the possibility that both our cognitive processes and our methods of evaluation could improve, and not just in degree but perhaps also in kind.

One possible way to make plausible the idea that we are doing the best possible job with our epistemic evaluations is to argue that they do not *merely* come to us through cultural inheritance, but that they do so with some special guarantee of their correctness. Alvin Plantinga sees this situation quite clearly and uses it to argue for the conclusion that if one is an epistemological naturalist (as he considers a proper functionalist to be), one must be a metaphysical supernaturalist. His argument goes something like this.[xviii] Suppose there is no divine being guaranteeing the proper function (or reliability) of our faculties. Suppose, that is, that our methods of belief acquisition evolved in an environment much like that which modern synthesis evolutionary biology suggests. In such an environment cognitive processes would be selected for their survival value and not necessarily their truth-conduciveness (and thereby their proper function and/or reliability). But having cognitive processes selected for their survival value does not guarantee that they will produce true beliefs. In fact, given the wide set of alternatives, there is a relatively low probability that they would produce true beliefs.[12] The

---

[12] Plantinga precisifies these arguments somewhat using Baysean theory, but I will ignore these details. For those interested, they can be found in the chapter already cited.

kicker, of course, is that because such improbably truth-producing cognitive faculties produced the theory of evolution, this theory is highly unlikely to be true and thus indirectly self-defeating. Plantinga's way out of this sticky situation is merely to suggest that the theist has a guarantee that her faculties are truth-producing, because she believes that God guided evolution to create man in God's image. Most epistemologists will want to avoid this conclusion for independent reasons, though. I would suggest that one could successfully avoid it by leaving externalism's epistemic conservatism behind. If one does this then one does not have to be worried that evolution is unlikely to have produced true beliefs because one is not necessarily concerned about which beliefs are true. Along with Stich, from whom Plantinga ironically borrows in making his argument for supernaturalism, one could merely be concerned about which beliefs have instrumental value.[xix]

The externalist may have other recourse due to her theory's residue of foundationalism. The argument might go something like this. Basic beliefs gain their warrant just by the fact that they come from cognitive faculties that we all trust—across times, cultures, languages, etc. They produce beliefs that we can all think are true or at least that we can all agree on. Mediately warranted beliefs gain their warrant by proper connection to basic beliefs. Unwarranted mediate beliefs arise due to improper connections with basic beliefs or connections with unwarranted basic beliefs. Since no issues with truth arise concerning basic beliefs and mediate beliefs depend on proper connection to basic beliefs (and not necessarily on their teleological relationship with truth), then problems with truth are avoided.

I doubt that this argument works. It seems that carving out properly basic beliefs relies on the same arbitrariness as is involved in an appeal to truth, but even supposing the argument goes through, the reliance on foundationalism is too much of a liability for externalism for this move

to be worth it. The problem lies with the proposed asymmetry of basic beliefs and mediate

beliefs. To see this we must examine how basic beliefs gain their warrant a bit more carefully.

Such an examination will problematize foundationalism beyond just the externalist version.

For the externalist foundationalist, it is not adequate to proceed just because we're not

violating any rule that we epistemologists have established (as seems to be the case for some

internalists); we must, in addition, have positive reason to think that basic beliefs are warranted.

Plantinga provides a positive reason indicative of the externalist viewpoint on basic beliefs (the

reliabilist picture would be slightly different, but, I think, prone to the same arguments that I

provide). According to him, a warranted belief is "produced by [properly functioning] cognitive

faculties according to a design plan successfully aimed at truth".[xx] We must be doing more than

just avoiding epistemic violations in accepting basic beliefs: we must be on the path to truth, and

we must be traveling that path with cognitive faculties that work correctly. That is not to say that

we must *know* whether or not we are oriented towards truth, but rather that the faculties we are

using for producing the beliefs must be faculties that are designed to produce true beliefs. So, for

instance, our beliefs should not arise out of wishful thinking or sycophancy or spite.[xxi]

If a belief comes to one immediately and has warrant (because it came to one through a

properly functioning cognitive faculty etc.), then it is said to be a properly basic belief. The

basic-ness is proper because one is trusting faculties that one should be trusting. But just because

properly basic beliefs are acquired with warrant does not render them indefeasible. If we are

confronted with good reason to drop it, then it is our epistemic duty to do so. Take the case in

which I believe that I saw Pedro on the streets of Omaha yesterday. This belief derives from my

memory of having seen Pedro[13], and as such is properly basic. However, if Pedro's trustworthy

---

[13] As Plantinga points out, this memory need not necessarily be a replay of the sensual experience of seeing Pedro, it could just be a 'doxastic experience' in which it just occurs to me that I did see Pedro yesterday.

boyfriend Jeff tells me that Pedro was traveling in Moscow yesterday and that his twin was in Omaha, I have a defeater for my belief that I saw Pedro yesterday. This does not mean that I was originally unwarranted in trusting my memory, but it does pull the warrant out from under my original belief so that it is no longer responsibly tenable.

Given their defeasibility, basic beliefs aren't basic in the way that the traditional foundationalist would conceive of them. Non-basic beliefs can defeat basic beliefs and vice versa. Even more damagingly, though non-basic beliefs by definition cannot be *evidential* supports of basic beliefs, they can support them. Specifically, which other beliefs I have can affect which basic beliefs count as properly basic. The state of my noetic structure determines how I form basic beliefs. If this is the case, it seems hard to hold onto externalism and foundationalism simultaneously, for then it is hard to see how basic beliefs hold a special status that avoids arbitrariness. To see that it is, let us examine examples.

Take the case where Alvin feels tired. He believes he is tired. His belief that he is tired is warranted because Alvin's proprioceptive capacities are in full working order and are telling him that he is tired. The relevant further question is: how does feeling tired bring about the belief that he is tired? It seems that in order to convert the feeling of tiredness into the belief that he is tired, Alvin must be able to interpret the feeling he has as tiredness. The ability to believe that he is tired based on his feeling of tiredness is, then, "an acquired causal characteristic" of Alvin, as Sellars has convincingly argued.[xxii] Nothing special about tiredness—save its familiarity—makes it a source of immediate justification. A very large set of beliefs can be properly basic if the agent undergoes the correct training. Alvin was trained to associate his feeling of tiredness with the term and concept "tired", but he could conceivably be retrained to believe that whenever he has the feeling he currently associates with tiredness, he is actually experiencing the pressure of a

magnet placed in his head. Suppose, for instance, his doctor informs him of this fact after a series of examinations. Then Alvin's belief that a magnet in his head has been activated would be immediately justified by the feeling that he once associated with tiredness. This is just a modification of Sellars's point that "observational knowledge of any particular fact...presupposes that one knows general facts of the form X is a reliable symptom of Y", which means that "one couldn't have observational knowledge of any fact unless one knew many other things as well".[xxiii] That is, the warrant of a basic belief depends on the warrant of many other beliefs.

But maybe only basic beliefs based on phenomenal experience work this way. Maybe the warrant of beliefs based on doxastic experience (see footnote 13 above) are independent of the warrant of other beliefs. This seems unlikely. Suppose Jeff asks Pedro what the first musical instrument he learned to play was. Pedro, whose memory is working perfectly well, recalls that he learned to play the oboe before any other instrument. The belief generated by this memory—that the first instrument he played was the oboe—depends on the warrant of many other beliefs of Pedro's. For instance, he must believe that an oboe is a musical instrument and that banging on pots and pans as a toddler does not count as learning to play a musical instrument. To see this, imagine if Pedro had lacked either of these beliefs before Jeff asked him about his creative genesis. If Pedro (unwarrantedly) believed that an oboe is not a musical instrument, then he would (unwarrantedly) believe that he never learned to play a musical instrument.

The important point here is that basic beliefs do not just arise from an experience through some occult process. They arise from an experience given that the believer has a series of other beliefs and has learned to associate certain experiences with certain beliefs. The association involved here, is, of course, implicit—the believer does not have an experience and then dig

through their catalog of beliefs to find the one that matches it—the belief is, rather, generated habitually. Plantinga's statement that the belief is "occasioned" by the experience remains true, but now the occasioning is not so mysterious. And once the occasioning of a basic belief loses its mystery, properly basic beliefs don't seem so special and privileged. Sellars's comment that "empirical knowledge…is rational not because it has a foundation but because it is a self-correcting enterprise which can put any claim in jeopardy, though not all at once"[xxiv] implies not only that we should not look to basic beliefs as the special part of our noetic structure that makes the whole rational[14], but also that basic beliefs should be put to the same scrutiny as any other. Which eliminates the asymmetry inherent to foundationalism,[xxv] and with it any potential recourse to this asymmetry in defense of externalism.

Without any special reason to think that philosophers' intuitive judgments about which of our beliefs are epistemically valuable have some special insight, the externalist's procedure has little traction to move forward. I have suggested various reasons to think that philosophers' intuitions do not have this special insight. If my arguments are successful, externalism is seriously lacking as an analysis of the concept of warrant. We cannot simply assume that we know what the good beliefs are and then analyze what they have in common to find out what underlies a warranted belief. We must be aware that when we attribute warrant to a belief, we are operating within some set of epistemic standards of value. In order to understand what goes into some given attribution of warrant, we have to understand what about the attributor (an externalist epistemologist, in this case) causes her to make such an attribution. Whereas externalists have taken themselves to have discovered what the *real standards* of warrant are, they have only

---

[14] …and not *just* rational. Plantinga examines what makes a belief justified, rational, and warranted, with warranted being the strongest claim about a belief. Because Sellars did not think in these terms, I think it's fair to say that we could replace 'rational' with 'justified' or 'warranted' or any related epistemic term of praise and still capture his intentions.

analyzed what the standards of warrant are for people like them in the situations that they have

investigated.

EVIDENTIALISM

Earl Conee and Richard Feldman's celebrated evidentialism is an updated version of internalism that remedies some of the problems with more traditional forms. The basic assertion of evidentialism is that a doxastic attitude towards a given proposition is warranted for a believer if and only if having that doxastic attitude towards the proposition *fits the evidence that the believer has at the time of the belief.*[xxvi] Conee and Feldman identify this position as internalism because all of the factors involved in the warrant of a belief are internal to the believer. Evidence, in their sense, can only be some mental state or series of mental states. However, Conee and Feldman insert an important caveat to their internalism. They note that someone who has evidence for something "need not know, or have good reason to believe" that they have that evidence.[xxvii] It merely must be the case that somewhere in the believer's cognitive system there exists appropriate evidence for her belief. Another aspect of evidentialism that distances it from traditional internalism is the evidentialist position that the believer's reasoning for such a belief is not the only thing that can count as evidence for it—any sort of mental state, from memory to perception to heard testimony, will do.

Due to these adjustments, some of the major difficulties for traditional internalism fail to stick to evidentialism. Whereas both traditional internalist positions had posited global properties of belief, evidentialism need not take any position on doxastic architecture or what special qualities of the whole render its parts warranted. One's belief must merely fit the evidence that one has for it—whether or not this evidence has global coherence or rests on a special set of beliefs need not come into play. It also avoids the internalist requirement that a believer be aware of the evidence for her belief in order for it to count as warranted. The believer must, in some

sense, possess the evidence in question (it cannot be evidence that exists somewhere unbeknownst to the believer),[15] but the believer need not have "some internal representation of this fact"[xxviii] to make that evidence relevant to the warrant of the belief. Just as importantly, evidentialism does not restrict the moniker "evidence" to just beliefs. Perceptions, memories, and testimony can all count as evidence for a belief. As long as the evidence is internal to the believer—as long as the believer at some point has *been aware* of the evidence—then it can count as a justifier/warrant-bestower for a belief. In short, evidentialism does not confuse warrant with justification.

Evidentialism also sidesteps some worries about externalism. There is no clear way in which evidentialism can be seen to rely on truth, with all its aforementioned difficulties. Indeed, if one has good evidence for a belief, one can be said to be warranted in holding it even if it is produced by an illusion.[16] The belief need not be produced by a design plan aimed at truth or a reliable belief-producing process; it only needs to have good evidence. Evidentialism can thus maintain the normative aspect of externalism without falling prey to the difficulties of saying what counts as true or a good design plan.

The problem with evidentialism arises when one asks what makes some bit of *potential* evidence count as *actual* evidence. Or, to put it in Feldman's terminology, what makes some bit of evidence *epistemically acceptable*?[xxix] Take the case of Polly, the paranoid schizophrenic. Conee and Feldman's position quite clearly implies that her paranoid belief that the local government has created a program to spy on her every move does not have sufficient evidence.[xxx] But suppose *she* thinks it has evidence. She has a series of mental states she can

---

[15] In the same volume, Richard Feldman argues for the view that the evidence must be present in the believer's mind at the time of the belief, but this is just one type of evidentialism. I will concern myself with the more general case that takes no position on this issue.

[16] Conee and Feldman's treatment of Putnam's brain-in-a-vat scenario seems to imply this conclusion (op. cit., 60-61)

refer to that she thinks provide evidence for her belief that the government is spying on her. Why does what she counts as evidence not count as evidence all told? We might say that she is not responsive to defeaters or that the amount of evidence she has is not really enough to warrant the belief that she is being spied on. Polly might think, though, that the government is just doing a good job covering its tracks—what seem like defeaters are really contrivances and what seems like insufficient evidence is just the most available given the cover-up.

In the evidentialist schema, we have a believer, a belief up for evaluation, and a series of other mental states possessed by the believer. If the belief in question counts as warranted, then the series of other mental states under consideration count as evidence for that belief. The question that arises in the case of Polly, though, is *when* these mental states should be considered evidence. If we don't answer this question, then it will be difficult to pick out which of our uncountable mental states should be considered in evaluating the evidential support for a belief. Yet Conee and Feldman never point to principled criteria for selecting mental states. They rely on our intuitions about how evidence and epistemic support work. The problem, though, is that not everybody has the same intuitions about what counts as evidence for what. Indeed, our judgments about what bits of potential evidence are epistemically acceptable may be situation-dependent.

At this point Conee and Feldman might refer to a notion that they call well-foundedness.[xxxi] For a belief of Polly's to be well-founded it must be warranted based on the body of evidence that she has and there must be no more inclusive body of evidence that she has with which her belief would better fit. Conee and Feldman might say, "let us grant that Polly's belief is warranted according to the body of evidence she has in mind. Even if this is the case, she has a more inclusive body of evidence that she *should* be using in evaluating the warrant of

her belief." One initial worry we might have about such a response revolves around the concept

of inclusiveness. What makes some body of evidence more inclusive than another? It would

seem hard to provide an answer to this question without reverting to the old-fashioned internalist

appeal to global features of beliefs. But let us leave that worry aside. Suppose, instead, that Polly

does not have a more inclusive body of evidence. In the case of a paranoid schizophrenic this is

not too hard to imagine: Polly may only have mental states that fit her paranoid belief. She may

never have entertained an alternative explanation for her experiences. It seems that her belief

would still count as well-founded according to the criteria set out by Conee and Feldman simply

because there is no more-inclusive body of evidence to appeal to in her case. What the

evidentialist might want to say here is that if Polly were in her right mind she would have a more

inclusive body of evidence. But making judgments about being in one's right mind also involve

normative judgments about which cognitive processes are good. There is, then, still an

unexplained appeal to normative principles at issue.

What comes out in examining the case of Polly is that we have normative intuitions about

what count as good and bad evidence and that other people may not share these intuitions. We

are thus in a similar place as we were in looking at the normativity implicit in externalism. In

figuring out what counts as acceptable evidence, one must take into account *who* is doing the

epistemic evaluation. Polly would evaluate evidence of the government's secretive surveillance

as acceptable evidence, while I would not. In any process of evidence evaluation, there is a

normative element implicit, and when a normative element comes into play we need to know

something about whose standards are being employed and why. Though the example at hand my

seem too absurd to matter, it seems clear that different normative positions on evidence exist

within those who think that religious experiences count as warranted and those who do not, and

that these differences are epistemically, if not practically, important.[17] Suppose Alvin just had a deeply moving experience that he interpreted as a direct experience of God. He recounts his experience to both Plantinga and Daniel Dennett. Plantinga will most likely interpret the content of the experience as evidence (for Alvin) for the existence of God, while Dennett will most likely deny that this experience should count as evidence for the existence of God. With whom one agrees depends on the details, but the point is that evaluations of evidence should not be seen as evaluator-independent. One must have a series of normative epistemic beliefs in order to engage in any such evaluations.

Conee and Feldman take for granted that the examples they provide will elicit certain epistemic evaluations. Like the externalist's appeal to truth, however, this agreement on evaluations may rely more on the shared cultural background of analytic epistemologists than any inherent quality of the beliefs and situations in question. It certainly relies on shared normative epistemic intuitions, and, as Stich has argued, we have no reason to think that these intuitions have any special value merely because they are shared. Research by him and others[xxxii] seems to indicate that our epistemic intuitions are not shared by all people from all cultures. Even apart from this research, we have no *a priori* reason to think that our evaluations are mostly reliable and much *a posteriori* reason to think that they are not.[xxxiii] What we count as evidence probably has more to do with contingent facts about how we were raised than necessary facts about the intrinsic value of certain beliefs or certain relations between beliefs and other mental states.[18]

---

[17] See Plantinga, *Warranted Christian Belief*. My response to Plantinga's position can be found in "Making Sense of the *Sensus Divinitatis*", unpublished.

[18] Williamson's knowledge-first approach might be considered an externalist form of evidentialism. For Williamson, the warrant of beliefs is based on the evidence of all and only one's knowledge. He understands knowledge as the basic factive mental state (every factive mental state is also knowledge) and counts something as evidence if it increases the probability of a given belief. Just as with Conee and Feldman, the evidential relations need not be

known by the subject in order to count as evidence. One does not even have to know that one knows something in order for it to count as knowledge.

Though I cannot discuss it in sufficient detail, it seems to me that Williamson's approach is beset with the difficulties of both evidentialism and externalism. Because knowledge is factive, in order for a belief to count as knowledge it must be true (this does not imply that knowledge is a property of a belief, it just provides a convenient way of talking about these issues). But again, how do we know if any given belief is true? Williamson might say that whether or not we know that a given belief is true does not matter, for we do not need to know that we know a given belief in order to know it *simpliciter*. There is a fact of the matter about whether or not we know something due to its objective relation with truth, even though we may not have any cognitive access to this fact of the matter. The warrant of a belief, then, depends on only things that one knows, whether or not one knows (or anyone else knows) which of one's beliefs actually count as knowledge. In other words, some of our beliefs have warrant and others do not, but we may or may not know which.

I think that the cost of such a thoroughgoing externalism outweighs the benefits. The big cost is that we don't (necessarily) know which of our beliefs are warranted, which leaves us open to skepticism. Williamson's responses to skepticism are not enough to close this opening. He claims that the skeptic begs the question by claiming that one cannot know if one is in a skeptical scenario in a non-skeptical scenario. The skeptic hopes to prove that we cannot know if we are in a skeptical scenario in a non-skeptical scenario, but she simply assumes that this is the case in going about proving it, or so Williamson thinks. He attempts to make the contrary palatable by arguing that our evidence is different in the real world than in the skeptical scenario. The problem is that the sort of evidence that Williamson has in mind is the type that we do not know if we have, so the skeptical challenge has merely been temporarily pushed back. Even if not necessarily knowing which of our beliefs are warranted does not leave us open to skepticism, it does put us in an unfortunate and puzzling situation. Recall that warrant is supposed to be what we should believe, so on Williamson's picture we are left saying that we do not know what we should believe.

Perhaps we should not, on reflection, have any problem with not knowing what we should believe. Then the difficulty then becomes how we should go about believing in our impoverished state. On a certain level our epistemology must take pragmatic concerns into account, for we do continue believing things and we want to be able to do so in the best way possible.

EPISTEMIC FORMS OF LIFE

Although there are more positions on warrant and more variations thereon, I take it that those I have addressed thus far cover most of the landscape. At the least, I take it that those I have addressed so far create a good schematic map of the general sorts of position that have currency. I have examined some specific problems with each position, and I would now like to suggest that these specific problems share some important characteristics. These problems must be remedied in order for a theory of warrant to account for and effectively guide our actual epistemic practices. Pragmatism about warrant provides this remedy by addressing their common faults.

Most of the time we can go about evaluating the warrant of a belief without taking into account our own role in the evaluation. We can do so because we are mostly surrounded (or perhaps mostly surround ourselves) with people that have similar epistemic intuitions as we do and with situations that do not test our evaluative limits. Taking advantage of this state of affairs, epistemologists have mostly gone about creating theories of warrant with the implicit assumption that there is one right way to go about evaluating it. Internalism relied on global structural features of a believer's doxastic architecture for standards of warrant, and the skeptic exploited this placement (and the exclusion of attributor factors) to provide an unanswerable challenge. Externalism located these standards in the more concrete realm of the functioning of our cognitive faculties, but left out the normative aspect of warrant attribution in favor of anchoring beliefs in truth. This approach failed because truth may not be as solid an anchor as it seems at first. Evidentialism moved the anchor from truth to one's evidence, but still failed to account for the normative standards employed.

As a proposed corrective to these views, I will examine why people value certain beliefs without assuming that everybody always does so for the same reason. To do this, I will propose a theory of how we come to share points of view on the warrant of a belief. It is this process, I argue, that makes a belief warranted. I call my view pragmatism about warrant because it relies on the pragmatics of human interaction to explain the workings of warrant. This reliance makes it similar to the views of the original pragmatists (especially William James), the neo-pragmatists Richard Rorty, Hilary Putnam, and Donald Davidson, and those more recent inheritors of this tradition like Michael Williams, John McDowell, and Stephen Stich. Like Rorty's, the pragmatism I develop below will see the evaluation of beliefs as the result of our practices and conversations, but, like Putnam's, it will allow for subject-independent standards for good and bad beliefs. In order to figure out what we should believe and why, I will engage mostly in the descriptive task of explaining how we determine the answers to these questions in our various discourses. In this, I shall be guided by Wittgenstein's principle that "what people accept as justification is shown by how they think and live".[xxxiv] I shall mainly defer prescriptive statements to the specific discourses themselves, but I will have some things to say about choosing between rival discourses and why some are better than others.

To get started let us examine a simple case of belief evaluation. 4-year-old Leila is looking out her window and says to her mother, "look! A squirrel!" Her mother looks out the same window and sees a squirrel. She judges Leila's belief—that there is a squirrel outside the window—to be warranted and praises her for it. In this instance of belief evaluation, and in many others like it, there are three key elements: a believer, a belief, and an evaluator. Leila, the believer, expresses (albeit somewhat indirectly) her belief that there is a squirrel outside of her window at that moment. Leila's mother, the evaluator, understands the belief that Leila has

expressed, and evaluates its warrant by looking out the window. In general, epistemologists have focused on the first two aspects of this triad: the believer and the belief. They have proposed aspects of the believer and the belief that affect the evaluation of the belief's warrant. By doing so, they have left out a crucial part of the picture: the evaluator. It shall be my concern to reinsert the evaluator into this picture and examine what it takes for her to make warrant evaluations.

The relationship between evaluator, believer, and beliefs can be usefully compared to that between interpreter, author, and text or to that between listener, speaker, and utterance. Using the former comparison, the work of Hans-Georg Gadamer can provide the initial steps to understanding the evaluation of warrant. Most of his work focused on textual hermeneutics, in which an interpreter interacts with a text and its author, but many of his more general considerations are analogous to those that concern us. Indeed, sometimes we perform epistemic evaluations on the beliefs stated or implied in a text. Aspects of his analysis can thus be co-opted for our purposes. One of Gadamer's most important insights is that "hermeneutics must start from the position that a person seeking to understand something *has a bond to the subject matter* and…has, or acquires, [this] connection from the tradition from which the text speaks"[xxxv]. When dealing with texts, this bond can have thick intermediary layers of tradition based on years of scholarly work, but even the act of interpreting some present subject's beliefs involves a shared connection to a subject matter (and a tradition). Both the evaluator and the believer must make sense of the beliefs in question, and doing so involves a shared way of thinking about the topic at hand and what sort of beliefs are appropriate in dealing with it. In engaging with the same subject matter as the subject, the evaluator participates in the same tradition—which will be roughly equivalent to what I will call "engaging in the same form of life"—as the subject.[xxxvi] While the evaluator's remove may allow her to reflect on the believer's situation and thus

evaluate the believer's beliefs more deeply than the believer herself may be able to,[xxxvii] the evaluator must have enough common ground with the believer to be able to make some sense of her beliefs. The act of making sense of a belief is the act of interpreting and we need common ground to interpret.

Donald Davidson has offered parallel considerations in comparing evaluator, believer, and beliefs to listener, speaker, and utterance.[xxxviii] He argues that in order to interpret another's statements, we must assume that we share a background way of thinking about the world with them. Without this background, we would be unable to make any sense of another's words. In order to share a background, we must take for granted that we see things in much the same way, or, to put it in his terms, we have a "general agreement on beliefs". This is not to say that we make some conscious decision about our agreement; rather, the very act of interpreting—of understanding what somebody is trying to say—automatically involves such an assumption of agreement. This agreement is closely parallel to Gadamer's concept of a bond to the subject matter. In both cases, and, by extension, in the epistemic case with which we are concerned, the item to be interpreted can only successfully be interpreted if the interpreter and the author/speaker have the same idea of what is being addressed by the interpreted item.

In the epistemic case, the evaluator and the believer must share *common ground* against which the belief can come into view. Leila's mother had to understand what Leila was talking about in order to even begin to evaluate the warrant of the belief that there is a squirrel outside the window. Epistemologists are in a similar position in evaluating the belief of any believer that comes up on epistemological discourse, even if that believer is a fictional figment of the epistemologist's imagination. I think a good first blush way to think about what goes into establishing the common ground between evaluator and believer can be found in the concept of a

form of life. Roughly, a form of life is a way one goes about a particular activity. When people

go about a particular activity in similar ways (which, of course, people generally will do if they

are actually engaged in the same activity), they share common ground at least on the way that

activity works and the subject with which it deals. When people share common ground on an

epistemic activity, an *epistemic context* arises in which the epistemic status of beliefs can be

considered. Thus, warrant evaluations are context relative. But let's not get ahead of ourselves.

First we must understand what exactly a form of life *is* and what it means for a believer and an

evaluator to share one. In order to do this we must first clear up the notion of common ground.

How we establish common ground has mainly concerned philosophers of language, but

many important theoretical constructs found in this work apply to forms of life more generally.

Specifically, work on coordination games following the example of David Lewis (who borrowed

the concept from Thomas Schelling) can help us understand how common ground gets

established in general. I will use as my guide the more recent work of Herbert Clark, who has

entered the conversation started by Lewis from the perspective of psycholinguistics and has

moved it forward both theoretically and experimentally.

Clark begins with a simple idea: "everything we do is rooted in information we have

about our surroundings, activities, perceptions, emotions, plans, interests. Everything we do

jointly with others is also rooted in this information, but only in that part we think they share

with us."[xxxix] This part is common ground. Common ground is a necessary condition for

successful joint activity. An activity, in his sense, is a "goal-defined, socially constituted,

bounded [event] with constraints on participants, setting, and so on, but above all on the kinds of

allowable contributions",[xl] and a joint activity is simply an activity where more than one

participant is involved. Warrant evaluation is a joint activity because it always involves at least

one evaluator and one believer.[19] Joint activities can be seen as solving the aforementioned coordination problems, where two or more people have "common interests, or goals, and each person's actions depend on the actions of the other".[xli] Such activities include playing a duet, rowing a boat, kissing, and playing soccer. In order to achieve the goals of these activities, one must both act successfully and coordinate one's actions with the other participants in the activity. The coordination with others is where common ground comes in. One must share common ground with the other participants to coordinate with them.

In order for a given item (belief, proposition, experience, etc.) to be common ground for a community, Clark argues (following Lewis), three things must be true of it.[xlii] (1) Every member of the community in question must have information that a given state of affairs holds. (2) The state of affairs must indicate to every member of the community that every member of that community has information that the state of affairs holds. (3) The state of affairs must indicate the item in question to every member of the community. Note that these conditions must apply simply for the members of communities to make sense of the item—in order for it to be meaningful for them. For a belief to be warranted, it must not only be interpretable/meaningful to the evaluator, it must also be evaluated positively by the standards she is employing. More on what causes this evaluation will come later—for now we are simply focusing on what it takes for a belief to be entered for consideration.

In any case, this tripartite formulation leaves the most crucial relation hanging in the balance: what does it take for a state of affairs to indicate something to somebody, and, more pertinently, how can a given state of affairs indicate the same thing to multiple people? One way

---

[19] The believer and the evaluator can be the same person. I can evaluate my own beliefs. However, I think in these cases one adopts a critical persona to evaluate one's own belief and thus simulates a joint activity. Clark argues that we might see talking to one's self as formatted similarly to a joint activity because conversation is the basic form of language use and we just mimic it in non-basic forms. I'm not sure that I want to endorse this controversial position, but it is a tangential issue that merits consideration elsewhere.

a given state of affairs can act as a basis for common ground in this way is for that state of affairs

to be part of a convention that the community shares. Lewis explores this possibility in rigorous

depth. I want to suggest that forms of life work roughly like convention does in these

circumstances, and in fact address more broadly what is needed to get at this relation.

The notion of a form of life is Wittgenstein's. In discussing how meaning works, he says

that "to imagine a language means to imagine a form of life".[xliii] And, a bit later on, "the term

language-*game* is meant to bring into prominence the fact that the *speaking* of language is part of

an activity, or a form of life".[xliv] I want to generalize the concept of form of life beyond mere

language use, and I think we can find the seeds for this generalization already present in

Wittgenstein's conception. As mentioned before, a form of life is a way that somebody goes

about a particular activity. By calling linguistic activity a series of language *games*, Wittgenstein

is bringing out the fact that language use is always embedded in an activity and that human

activity is always governed by some set of rules. In joint activities more generally, people

proceed according to rules—explicit or implicit. Though many joint activities use language to

accomplish their goals (to win their coordination games), not all do.

Different forms of life have different sets of rules. Physicists are physicists because they

engage with a certain type of information in a certain way—whatever makes this type of data and

this method of engagement distinctive constitutes the form of life of a physicist *qua* being a

physicist (i.e. when dealing with the subject of physics). The physicist's way of going about

things allows her to participate in contexts where different aspects of physics are under

discussion or contemplation. One could similarly find forms of life for mathematicians, spiritual

healers, epistemologists, jazz musicians, or any other method of meaningful (that is,

semiological) engagement with the world. Any sort of joint activity brings with it a form of life and thus a set of rules for engagement with that activity.

But what does all of this have to do with epistemology? Wittgenstein regarded the certainty involved with knowing "not as something akin go hastiness or superficiality, but as a form of life", though he thought this was "very badly expressed and probably badly thought as well".[xlv] What I think is badly thought about this aphorism is the notion that we have to deal with *certainty* in dealing with the value of beliefs. What Wittgenstein wanted to get at by addressing the certainty involved with knowledge (in using the concept of certainty he was merely following the Zeitgeist, specifically, the precedent set by Moore), was the quality associated with beliefs that acted as justifiers for other beliefs without themselves being justified. More revealing in this respect is his comment that the end of justification "is not an ungrounded presupposition: it is an ungrounded way of acting".[xlvi] This ungrounded way of acting need not involve certainty or even knowledge. It need only involve the simple fact that what an evaluator take as good reasons for some belief depends on the context in which that belief arises, which depends on the common ground between the evaluator and the believer, which depends on which form of life the evaluator and the believer are participating in. A form of life brings with it rules for engagement, which includes common ground on standards for warrant within that form of life.

Relatedly, Martin Heidegger thought that knowledge was a "founded mode of being", which means, roughly, that the standards for knowledge arise out of prerational action in the world.[xlvii] This insight can be adjusted for our considerations of warrant. In our dealings with the world, we gain implicit understanding of it. This understanding is beholden to our practical engagement—how we see the world depends on what activities we are engaged in within the world. Edmund Husserl reframed Heidegger's notion in this way: "for the human being in his

surrounding world there are many types of praxis, and among them is this peculiar and historically late one, theoretical praxis".[xlviii] Theoretical praxis (conceptual interaction with the world) maintains certain important qualities of more primitive manners of engaging in the world.

Less developed manners of interacting with the world differ depending the material with which they deal and the outcome that they hope to accomplish (think of the differences between, e.g., carpentry and husbandry), and conceptual interaction with the world maintains this pattern. Indeed, as Clark, Wittgenstein, Heidegger, and Husserl have all pointed out, conceptual interaction with the world works *only in service of* these less developed activities. The differences between activities underlie the varying standards for judging human behavior—one can only judge the quality of the actions within a given activity based on the pragmatic considerations of that activity. Warrant is one type of standard that we use for certain types of activities. To put it another way, given the fact that warrant arises in a founded mode of being and the fact that our activities generally have at least a vague *telos* that guides how they proceed, we can conclude that standards of warrant depend on something like the *telos* of a given set of behaviors. If one does not like the concept of built-in teleologies, then the same point can be expressed by specifying that standards of warrant depend on what the activity at hand sees itself as doing (or is seen as doing when viewed correctly). I think, though, that the concept of teleologies employed here should be somewhat uncontroversial. It only implies the goal-directed nature of activities already briefly discussed above in the context of Clark's views. If Clark's arguments for conceiving of joint activities as solutions to coordination problems are successful, then it takes only the short step just mentioned to see that standards of warrant depend on the *telos* of the activity in question/coordination problem to be solved. But even if the concept of

teleology still seems problematic, the salient point is that the methodology of a given form of life dictates how warrant can be accomplished within that form of life.[xlix]

There is a slightly different way to see how forms of life create contexts for warrant evaluation. Engaging in a form of life excludes the discussion or questioning of a certain set of beliefs—the beliefs that form what Wittgenstein would call the "hinge" of a discursive form of life. Michael Williams has argued that this exclusion depends on the "methodological necessity" implicit in any given form of life. Given the methodology of the discourse at hand (the rules of the activity), we cannot move forward with the discussion unless we simply leave aside certain questions and lines of argument.[l] As Sellars has argued, our discourse "is a self-correcting enterprise which can put *any* claim into jeopardy, though not *all* at once".[li] By setting aside certain questions and lines of argument, the evaluator and the believer delimit the questions and lines of argument relevant to the present context. Their very engagement in a given subject matter generates methodological necessity. When two or more human beings share a form of life, they agree on a set of rules for engagement with each other and with the subject of discussion. These rules make some subjects of discussion off-limits in the sense that to take on such subjects would be to adopt a distinct set of rules;[lii] it would be to take on a different form of life. People do not just agree on a set of rules because their opinions on the subject happen to overlap, however. As with Wittgenstein's rules for language use, the agreement here "is not agreement in opinions but in form of life".[liii] Evaluator and believer are bound by the subject matter, by their shared background in the topic at hand. The very way they go about their discourse establishes its subject matter, which beliefs go unquestioned while engaging in it, and its standards of warrant. "It is anchored in all [their] questions and answers, so anchored that [they] cannot touch it". [liv]

Because the way different groups of human beings interact with the world together is often glossed as "culture", we might think of form of life as the aspects of culture relevant to epistemology. Or perhaps I should simply say that forms of life are epistemic culture. I mean culture here in a fairly broad sense, such as that defined by psychologist Triandis as "a pattern of shared attitudes, beliefs, categorizations, self-definitions, norms, role-definitions, and values…organized around a theme"[lv] or, especially, by psychologists Markus et al. as "interpretive frameworks".[lvi] "Culture" in this sense includes everything from the contemporary western world to medieval Muslim theologians to trekkies. As culture is notoriously difficult to define, searching for a clear-cut definition of form of life may be quixotic. Nevertheless, a sketchy outline of categories of things shared when people share an epistemic culture/form of life would include:[lvii] language, a set of experiences, acquired knowledge and/or prejudices (i.e. common sense), logical intuitions, notions of power/authority, et al. Through these categories and others, forms of life delimit epistemic context, or the specific state of some discourse.

As with culture in general, we gain forms of life mostly through inheritance. We use our cultural inheritance as common ground to engage with our cultural peers; indeed, the common ground itself makes possible such engagement. If we think of forms of life as roughly equivalent to what McDowell calls "second nature", then we can turn to his analysis (and that of Gadamer, from which he borrowed) to understand more deeply how epistemic inheritance works. McDowell suggests that we "generalize the way Aristotle conceives of the molding of ethical character [to] arrive at the notion of having one's eyes opened to reasons at large by acquiring a second nature".[lviii] He follows the German tradition in calling this cultural awakening "Bildung". Through the process of Bildung, we learn the demands of rationality, "which are there in any case, whether or not we are responsive to them"[lix] (this point is essential to understanding

Bildung, and will be developed further below). Bildung transmits a manner of looking at the world, a series of prejudices (good or bad),[20] that guide our epistemic lives. Wittgenstein argues that in order to follow a rule, we must have been trained to do so. A rule only guides "insofar as there exists a regular use of [the rule], a custom".[lx] Bildung passes such customs down to us. Through Bildung, we learn what sorts of questions to ask, how to answer different sorts of questions, and (at least roughly) why to focus on these questions.[21] But also as with culture in general, we can and do question any aspect of our inheritance. We can only question a given aspect, however, if we use others as friction. It remains impossible to interrogate from outside any context whatsoever. The best we can do is to question one form of life from the perspective of another.

Postulating that warrant is relative to contexts created by cultures may smack of the more controversial conclusions of Rorty. Specifically, it may sound that I am saying that a warranted belief is just "what our contemporaries will let us get away with saying".[lxi] Despite the fact that I think Rorty is often taken out of context, especially in his more flamboyant rhetorical flourishes, I want to distance the pragmatism I propose here from this (perceived) position of Rorty's rather than engage in an exegetical dispute on his views. Though in a literal sense I *am* arguing that warrant is what our contemporaries will let us get away with saying and this position *does* imply relativism, the specifics of how these sound bites are cashed out should make my position less absurd-seeming. Along with Stich, I think that relativism often gets short shrift due to cursory considerations of its details.[lxii]

---

[20] See Gadamer on the possibility of good prejudices in *Truth and Method*, 271-273 et passim.

[21] See Heidegger on *das Man* and authenticity for an interesting explanation of how culture justifies itself, albeit incompletely (*Being and Time*, passim) and Gadamer on the importance of questions (Truth and Method, Section 2, Chapter 4, Section 3, Part C). I might also note the parallels here to the Reidian position articulated by Plantinga on the basic-ness of testimony (cf. Plantinga, *Warrant and Proper Function*, Ch. 5).

When we evaluate one another's statements, we do so based on non-arbitrary standards of evaluation. We are governed by the forms of life that underlie our language games. These forms of life are unified by standards of grammar in the Wittgensteinian sense, which is a series of rules that we must follow in order to communicate effectively at all. Along with Wittgenstein, I "[do] not deny that that there is an independent or material reality that we can interact with; the claim is that this reality has no semantic content outside of language—that we understand it as reality only as it is acted out through semantic practice".[lxiii] Thus, "grammar tells us what kind of object anything is"[lxiv] because our forms of life delimit the subject matter of whatever we talk (or think conceptually) about. Grammar does not *make* the objects what they are, it only allows us to think and speak about them and as such gives them the semantic content necessary to be parts of statements that can be candidates for warrant.

The use of the words "grammar" and "rule" implies that the way we go about a discourse is not just whatever we can get away with—in the sense that our mistakes that go under the radar don't count as mistakes. We are guided by real standards, but, as John McDowell cautions, these real standards should not be given special ontological status.[lxv] We *can* fail to meet the standards of a given discourse, but the standards of that discourse *do* depend on the associated activity for their existence. It is notoriously difficult to explain this insight, but one finds it across a range of 20th century philosophers. Hilary Putnam, for instance: "What is better and what is worse to say about most questions of real human concern is not just a matter of opinion. Recognizing that this is so is the essential price of admission to the community of sanity".[lxvi] Notice that recognizing that this is so is not a matter of grasping some independent and transcendental standard; it is a standard implicit and immanent in our classification of statements and people as sane or insane. Gadamer, who worked in an entirely different tradition, had it that "[guiding principles] are not

norms to be found in the stars, nor do they have an unchanging place in a natural moral universe...nor are they mere conventions, but really do correspond to the nature of thing—except that the latter is always itself determined in each case by the use the moral [or rational] consciousness makes of them".[lxvii] Again, to the extent that our normative epistemic principles correspond to anything, they correspond to how we understand the world, to the traditions that we have inherited. The following of rules is just part of what it is to successfully do any sort of joint activity. Wittgenstein cautions that "to *think* one is obeying a rule is not to obey a rule".[lxviii] Here I think Crispin Wright's interpretation of Wittgenstein is helpful: "the grasp of [rules] is not anterior to the ability to give them competent linguistic expression but rather *resides in that very ability*".[lxix]

So discourses do have standards, but not standards underscored by a transcendent principle for warrant. For many, this may seem a halfway solution. After all, either our cognitive faculties are in touch with reality or they're not; either our reasoning is good or it isn't—regardless of what we think about it. Concerning the former, whether or not our faculties are in touch with reality depends on what we mean by reality and the particular metaphysical theory of what it means to be in touch with it. These are questions beyond the scope of a humble essay on warrant. Concerning the latter, we must realize that we do have standards for reasoning that can be misapplied, and hence on some level do not depend on what we think about them. This should become clearer in the forthcoming, and I will address it explicitly in more detail later on.

Now that we have a clearer picture of what a form of life is, let us see how it provides the appropriate relation for the three conditions for something to be common ground proposed by Lewis and Clark (who pioneered this territory). We should consider them not in their general

form but in their epistemic manifestations; that is, we are analyzing what it takes for a *belief* to

be common ground. (1) The evaluator and the believer must have information that they are in a

certain epistemic context. Evaluator and believer have this information simply by virtue of

participating in the same form of life. Physicists know that they are talking about physics simply

by engaging in the behaviors typical of a physics discussion. Many aspects of the form of life of

physicists could indicate the context: the fact that one is reading a physics journal, the fact that

one is in a physics classroom, the fact that one is speaking to a physicist, or even just the fact that

one is discussing topics that physics addresses. (2) The epistemic context must indicate to the

believer and the evaluator that they both have information that they are in that context. Because

the form of life provides each with information that they are in the context and any given form of

life is just a type of joint activity, the form of life should also indicate to each that the other has

similar information. In other words, their source of information is a shared means of dealing with

the world, so their information should also be shared and this fact should be transparent. (3) The

context must indicate some belief to both the evaluator and the believer. For those with

structuralist or holist tendencies, the way beliefs get set off by forms of life can be thought of in

terms of the Saussurean insight that the meanings of words arise from their contrast with other

words,[lxx] but, for those who find holism distasteful, it can also be thought of in less theoretically-

laden terms. Suppose Willard is telling Wilfrid why El Greco is his favorite Spanish painter and

he mentions that he especially admires his masterly work in *Las Meninas*. In order for Wilfrid to

evaluate Willard's implicit belief that El Greco painted *Las Meninas*, Wilfrid must have some

idea of who El Greco is, what *Las Meninas* is, who the painter of *Las Meninas* was, what

painting is, etc. He must have some idea of these things to even understand the belief that

Willard is communicating, let alone evaluate its warrant. If Wilfrid does not have some

knowledge of these things, then Willard's belief may as well be gibberish or he may as well not

be expressing any belief at all.

THE INSTABILITY OF CONTEXT

So far I have been writing as if forms of life can be clearly delimited from each other. I noted obviously established forms of life that have relatively evident and unchanging sets of standards implicit in them. One might even accuse me of having reified forms of life. It eased my exposition of the basic concept to act as if things were so cut and dried. But while it was convenient to talk about clear distinctions between forms of life, it was not entirely accurate. At least three considerations problematize my simplifications. First, forms of life change over time. Plant biology today differs strongly from its predecessors in the preceding centuries. Epistemological discussion of the nature of knowledge has seen various dramatic shifts in the even shorter time frame of the latter half of the twentieth century. Along with such changes come alterations in rules, standards, and even subject matter. Second, forms of life are not separated by clear dividing lines. A chemist may dabble in physics or biology—can we set off clear boundaries between these disciplines? Is a chemist not doing chemistry when talking about thermodynamics? When one uses a religious justification for a political belief, are they talking about religion or politics? We could create arbitrary cut-offs, but that would not establish principled differences between forms of life. Third, new forms of life can be created. Experimental psychology branched off from philosophy and gradually developed its own methodologies. Logic merged philosophy and mathematics together at around the same time. Evolutionary biology only came into existence after Darwin's work.

Feyerabend notes that, unfortunately for theoretical simplicity and for strict standards of warrant, "cultures change, they interact with other cultures and the indefiniteness resulting

therefrom is reflected in their [changing contexts]".[lxxi] In terms of Wittgenstein's extended game metaphor for forms of life:

> We can easily imagine people amusing themselves in a field by playing with a ball so as to start various existing games, but playing many without finishing them and in between throwing the ball aimlessly into the air, chasing one another with the ball and bombarding one another for a joke and so on. And now someone says: the whole time they are playing a ball game and following definite rules at every throw. And is there not also the case where we play and make up the rules as we go along? And there is even one where we alter them as we go along.[lxxii]

Even though our activities in the world are not cleanly delimited from each other and change along with changing interests, fashions, and observations, they still have rules that guide them. However, the indefiniteness of shifting forms of life does affect how rules work. When we consider forms of life as more amorphous entities, we should see rules as working *horizonally*. The concept of a horizon comes from Edmund Husserl. He posited that as we think about the world, we are "forever surrounded by an atmosphere of mute, concealed, but cofunctioning validities", which form a "vital horizon".[lxxiii] (Note the parallel to Wittgenstein, who thought that one's hinge propositions were so embedded in one's activities that one "could not touch" them.)[lxxiv] These validities surround us because of how we engage with the world, because of our form of life. They are horizonal because they point to a series of other validities that go on indefinitely, just as a literal horizon goes on indefinitely by pointing to other parts of the landscape. As one engages in different forms of life (however vaguely delimited), which validities are most relevant, and thus in the clearest focus, shifts. Jürgen Habermas, taking up Husserl's notion, has it that horizons "shift with the theme" of the activity at hand, just as "the

[literal] horizon shifts according to one's position".[lxxv] As we fluidly move from one form of life to another (and as we create new forms of life), the rules for action change. To say that rules are horizonal is to say that common ground is established and modified *en passant*.

Let us examine how this works in a concrete believer-and-evaluator case.[22] Suppose Edmund and Jürgen are having a conversation. Edmund expresses to Jürgen his belief that that nuclear war will wipe out mankind before global warming can even come close. Jürgen thinks that Edmund should not hold this belief, and he can appeal to various reasons why not: unchecked global warming will act faster than Edmund thinks, Edmund has not considered the factor of mutually assured destruction, Edmund should think that Obama's nuclear nonproliferation treaties are steps in the right direction, etc. Which of the reasons that Jürgen appeals to will shift the topic of discussion slightly, and along with it the rules and standards by which Jürgen and Edmund can evaluate each other's beliefs. If Jürgen appeals to his faith in Obama's actions, then Edmund would judge his belief based on international political grounds. Which facts and analyses the two of them will appeal to in such a discussion would be quite different than if they had taken up the discussion of how fast global warming will take effect.

No matter which of the discussions the two take up, they will have to take for granted certain beliefs and rules in order to move forward with it. These taken-for-granteds serve as the common ground that allows them to take part in the joint activity of conversation. Within a given conversational context, a form of life "appears as a reservoir of taken-for-granteds, of unshaken convictions that participants in communication draw upon" while engaging in discourse.[lxxvi] Notice that these convictions are *unshaken* and not *unshakeable*. We take certain things for granted because we must in order to engage in any sort of meaningful discussion—which is to

---

[22] In this conversation, and in many others, the roles of evaluator and believer switch back and forth depending on how the conversation goes. At any given statement of a belief, however, one could pick out a believer and an evaluator.

say, any discussion at all—but we could, in principle, challenge any of these taken-for-granteds. To do so, however, would be to reorient the focus of discussion and shift contexts. Recall, after all, that taken-for-granteds have a horizonal structure. What we take for granted in one situation could be under the highest level of scrutiny in another, but other beliefs and rules would have to be taken for granted in that other situation.

But how do interlocutors know that they are taking the same things for granted? How is common ground established without clear delineations between forms of life? More pressingly for epistemology, how does an evaluator know that she and the believer are following the same rules and taking the same beliefs for granted? To answer these questions, we must reexamine the interpretation relation that I discussed when first introducing the triad of evaluator, believer, and belief. In doing so, we will notice two seemingly incompatible tendencies. On the one hand, the evaluator must use empathy to try to figure out what the believer has in mind. In conversation, this will mean that the evaluator and believer must cooperate to establish common ground. On the other, in order to understand what the believer has in mind, the evaluator must make use of a bit of egocentrism. To see the believer as rational, the evaluator will have to compare the believer to herself.

What I have in mind here is a variant on the Davidsonian principle of charity. To begin my development of the concept, I want to start with Lewis's version from "Scorekeeping in a Language Game". [lxxvii] Once again, Lewis effectively uses the technique of comparing everyday human activity to a game: in this case to the specific game of baseball. In baseball, the score of the game at any given time can be given by a septuple of numbers representing balls, strikes, inning, runs, etc. The numbers on this theoretical scoreboard determine what plays are permissible by players in service of their goal of winning the game (which involves keeping

certain numbers of the score high and certain numbers low). Because the score at any given time depends on what has happened in the game up to that time and which actions are permissible at a time depends on the score at that time, which actions are permissible at a given time depends on what has happened in the game up to that time.[lxxviii]

Lewis then compares the game of baseball to a language game, which we can borrow for our different sort of joint activity: a justification game, we might say. In fact, Lewis himself explicitly cites the possibility of warranted assertion as a possible component in scorekeeping.[lxxix] The parallel is as follows:

Like components of a baseball score, the components of a conversational

score at a given stage are abstract entities. They may not be numbers, but

they are other set-theoretic constructs: sets of presupposed propositions,

boundaries between permissible and impermissible courses of action, or the

like. What play is correct depends on the score. Sentences depend for their

truth value, or for their acceptability in other respects [including warranted

assertion], on the components of conversational score at the stage of

conversation when they are uttered.[lxxx]

The score of the conversation, furthermore, "evolves in a more or less rule-governed way".[lxxxi]

What Lewis is offering here is a way to see the establishment of common ground as governed by forms of life while allowing the rules implicit in forms of life to be horizontal. Rules become strongly situationally dependent: which sort of beliefs will count as warranted depends on the state of discourse at a particular time. One of the key reasons this analysis is effective is because it breaks score into a series of components. Rather than seeing preestablished and rigid forms of life as establishing the common ground necessary to make epistemically evaluative judgments, we should view common ground as created by various factors governed by their own

less rigid forms of life. But what would these factors be? I have already suggested some possibilities above: language, a set of experiences, acquired knowledge and/or prejudices (i.e. common sense), logical intuitions, and notions of power/authority. Clark cites a different set of possibilities: community identification, shared expertise, communal lexicons, assumptions about human nature, norms about how conversations go, present environment (shared perceptual information), et al.[lxxxii] Triandis points to a related set: shared attitudes, beliefs, categorizations, self-definitions, norms, role-definitions, and values.[lxxxiii] Though I think there is some use in proposing which factors are at work in our creation of common ground, I don't want to favor one componential analysis over another. What I want to do, rather, is see what these componential analyses have in common and what that tells us about our epistemic behavior.

In order to take hold, components of common ground must be shared by evaluator and believer. Lewis thinks that both believer and evaluator will have a representation of conversational score at any given point in the conversation.[lxxxiv] This representation, I think, should not be taken as some conscious assessment of the situation on the part of the participants. Rather, which moves participants deem appropriate will indicate (and depend on) their implicit understanding of the state of discourse. If believer and evaluator have two very different score representations, then they will be unable to effectively communicate and the evaluator will be unable to interpret the believer's belief. However, participants rarely, if ever, *explicitly list* their interpretations of the present state of discourse. Generally, the evaluator will have to *interpret* the believer's understanding of the state of discourse by observing the actions the believer takes and inferring based on general patterns of score-interpretation to behavior pairings.

In order to interpret these pairings as fitting a given pattern, the evaluator's form of life must overlap with the believer's. These overlaps can occur on multiple levels of simultaneously

increasing scope and decreasing specificity. For instance, in order to share a language, believer and evaluator could share anything from grunts and gestures that somebody from most any culture could understand to English that any English speaker could understand to the dialect of English that speakers from Savannah, Georgia speak. In order to share norms, one might share anything from what the difference between a moral and immoral act is (not a specific theory on what differentiates them, just the ability to differentiate them) to specific views about the wrongness of free trade zones in third world countries. These levels can be seen as nested within each other: to be a Baptist is to be a Christian; to know about the ecology of a given ecosystem is to know things about population biology.[lxxxv] Sharing higher-level (more specific and shared by fewer people) factors of common ground, then, implies sharing lower-level components of those factors. We cannot speak the jargon of French art criticism without speaking French. We cannot talk about the merits of a particular baseball player without implicitly agreeing on what a baseball player is.

Suppose John says to Stephen: "I think Ryan Howard will break the record for home runs in a career". Stephen's ability to interpret John's statement depends on which components of common ground he shares with John and at what levels they share those components. For instance, Stephen might know very little about baseball, but he knows that home runs happen in baseball. Thus he could determine that Ryan Howard is somebody who makes home runs (whatever that means—something baseball-y) and that John thinks that he is very good at making home runs. He could also determine that John knows (or purports to know) a good deal about baseball, that he speaks English, and that John thinks that there is a record for home runs. Stephen would not be able to do a very good job evaluating the warrant of John's belief, though. He does not share enough taken-for-granteds with John to successfully move forward with a

discussion of John's assertion. His representation of the conversational score is different enough from John's that belief evaluation could not move forward successfully. If, on the other hand, Stephen knew a great deal about baseball, he and John would have similar representations of conversational score (even if they did not agree on the particular assertion at hand) and thus be able to take for granted a great deal more in moving forward. Stephen will be able to evaluate moves that John makes as appropriate or not because Stephen can characterize John's representation of score well.

It still remains to be seen how the evaluator characterizes the believer's representation of score, which leaves open the prior question of how an evaluator knows that the believer is following the same rules and taking the same beliefs for granted. To see how this works it will help to explore what how interpretation works when anomalies occur—when a participant has seemed to break a rule—and then to generalize aspects of these behaviors to all interpretation. Lewis argues that in anomalous cases, participants generally follow the principle of accommodation. The principle of accommodation states that whatever action is done should be interpreted as having been done according to the rules. In other words, participants will adjust the conversational score to accommodate for acts that do not fit with the previous score.[lxxxvi]

As I mentioned before, the notion of an accommodating tendency comes from Davidson (and, before him, Quine), who calls it the principle of charity. In Davidson's formulation, in order to understand others, "we must count them as right in most matters".[lxxxvii] So, in cases where a seemingly absurd belief is espoused, we generally accommodate it by assuming that, for instance, the believer uses her terms slightly differently or has uttered a malapropism, rather than thinking that the believer is woefully out of touch with reality. We do this because "if we merely know that someone holds a certain sentence to be true, we know neither what he means by the

sentence nor what belief his holding it true represents", and as such must interpret both the meaning of the sentence and the belief that he holds true.[lxxxviii] We cannot perform this interpretive action without assuming a general agreement on beliefs, because the possibility of interpreting another's belief successfully depends on taking similar things for granted. It is hard for us to even know what to make of somebody who asserts a belief that we think outrageously mistaken. Generally it will make more sense to interpret such a believer as meaning something different from what we mean rather than to think she could believe something like *that*.

The principle of charity comes out especially clearly in anomalous cases, but Davidson argues that it applies to all cases of interpretation. In general, we understand others' beliefs by considering them generally warranted by our own standards of warrant. Stephen Stich refines this point significantly. He argues that in order to have an *intentional characterization* of a believer's beliefs at all, an evaluator must adopt something like the principle of charity. Following Richard Grandy, though, he modifies the principle of charity into the principle of humanity. The principle of humanity has it that when we interpret "the pattern of relations among beliefs, desires, and the world", we must interpret them to be "as similar to our own as possible".[lxxxix] For most cases the principle of charity and the principle of humanity will overlap, but in some they will not. Suppose Taylor walks into the library and sees somebody that looks like Kanye sitting at a desk. She thinks to herself "Oh no! Kanye is sitting at a desk in this library!" However, the man sitting at the desk is not Kanye, but, coincidentally, Kanye is sitting at another desk in the same library. The principle of charity seems to instruct us to interpret Taylor's belief as true, but it seems that in fact she holds a false belief about the man that *looks* like Kanye. The principle of humanity recommends that we interpret his belief as false, because, in the same situation, our belief would be about the man we saw (and thus false), not the man we

had no thought about. In order to successfully characterize another's beliefs, we must see them as "reasonably similar to our own". In the case of a conditional belief, "if a [believer's] mental state does not interact with other mental states in a pattern which approximates the pattern exhibited by our own conditional beliefs, *it does not count as a conditional belief*".[xc] A similar argument would go for perceptual beliefs, beliefs about baseball, etc. What it means for a mental state to be a belief of a certain sort is for that mental state to behave like what we think beliefs of a certain sort behave like (what our own beliefs of that sort behave—or would behave—like). If we were to encounter a mental state that behaved very dissimilarly to any of our beliefs, we would have a hard time characterizing it as a belief.[23]

It is important to be clear about what it means for a mental state to be *similar enough* to our own to count as a belief. On the one hand, we obviously think that people that disagree with us on a large number of things have beliefs. On the other, there are some purported beliefs that seem so outlandish that we have trouble assimilating them into our concept of belief; we might think that the person who seems to hold them is in fact being disingenuous or doesn't properly understand the issue or is using different terminology. To get at this happy medium, I think we should turn back to Lewis's concept of the principle of accommodation. By asking how dissimilar somebody's mental states need to be for an evaluator to struggle in characterizing them as beliefs, we are asking how many rules a believer needs to break in order for the evaluator to be unable to accommodate their "play" (their mental state, their utterance). What goes on when an evaluator compares a believer's mental architecture to hers is not as egocentric or chauvinistic as it may sound at first—or as Stich accuses it of being.

---

[23] I'm not sure that "mental state" is the right term to be used here, because I'm not sure that in interpreting somebody's behaviors and words as indicative of this or that belief we are always pointing to mental states of that person. However, I will follow Stich's usage here for ease of exposition.

To interpret a believer's mental state, an evaluator must share common ground with the believer. To establish common ground with a believer, an evaluator must match up her own forms of life (components of conversational score) with those of the believer. If a believer seems to believe something anomalous to the rules of the forms of life involved in a particular context, then the evaluator has to adjust her representation of the context (the conversational score) to accommodate these beliefs as best she can. Interpretation requires overlaps in forms of life, and thus to accommodate the evaluator must assume that the believer is following the rules of a form of life that they share but that is slightly different from that of the context in which the evaluator saw them as being in. Even in cases where accommodation is not obviously necessary, evaluators interpret believer's behaviors and utterances as according to some form of life with which the evaluator is familiar. Accommodation is often tacitly involved in this interpretation due to the horizonal nature of common ground. If this form of life involves beliefs, warrant, and the like, then the evaluator will interpret behaviors and utterances as expressing beliefs that may or may not have warrant.

One could go into much more detail over what it means for something to be a belief. However, this essay is concerned with the next stage of belief evaluation—what it takes for a belief to be warranted. The evaluation part. The basic concept of warrant evaluation is simple: in order for an evaluator to interpret a believer's belief as warranted, the evaluator must think that the believer has followed the rules of the forms of life that have formed the context for that belief. To extend the analogy of justification games: the epistemic status of a belief depends on the conversational score when it is uttered. Because the evaluator cannot interpret the behaviors (linguistic or otherwise) of a believer as a belief without interpreting that belief as embedded in a specific context—as attempting to follow the rules of some form of life—the evaluator cannot

judge the warrant of a belief without asking whether the believer has followed the rules the form of life in which the behavior occurred.

Recall, though, that to create any given context for warrant evaluation, multiple components of common ground on multiple levels are being employed. To see this concretely, let us return to the case of John and Stephen. When John enters his belief that Ryan Howard will break the record for career home runs, Stephen's ability to evaluate the warrant of this belief depends on whether or not he shares enough components of common ground with John at a high enough level. The Stephen who does not know about baseball will be quite unable to evaluate the warrant of John's belief because, though he shares some forms of life with John at varying levels, he does not share the form of life of a baseball fan with him at even close to a high enough level. He cannot take for granted enough beliefs about baseball to make such a judgment. However, if John had entered the belief that Tom Cruise will break the record for career home runs, Stephen would likely be able to evaluate this belief even without much knowledge of baseball at all. Because Stephen knows that Tom Cruise is an actor and not a professional baseball player, he will be able to tell that John has made a mistake (or a joke) in entering his belief that Tom Cruise will achieve some accomplishment that only professional baseball players can achieve. In this scenario, Stephen needs only his low level knowledge of baseball to be able to judge the warrant of John's belief. He does not have to take much knowledge of baseball for granted.

What evaluators evaluate in deeming a belief warranted or not is roughly what Michael Williams calls a believer's *epistemic responsibility*. The more epistemically responsible a believer is, the higher *personal justification* she has in holding a belief.[xci] A person is personally justified if she acts according to the rules by which the evaluator is judging her belief. However,

epistemic responsibility goes both ways. An evaluator must be epistemically responsible by maintaining the common ground on which to judge a believer's belief. This basically means that an evaluator should abide by the principle of humanity, with the added notion that an evaluator should not just judge whether or not a behavior counts a belief, but also that a believer is acting how an epistemically responsible person would act. Using the principle of humanity, this latter judgment becomes a question of how the evaluator would act in a similar epistemic context.[24] Generally doing so will mean interpreting the epistemic context charitably for the believer. That is to say, the evaluator will mostly interpret a believer's belief as having been entered in an epistemic context in which it is warranted. Sometimes, of course, this will be impossible. In some cases this will involve failures of intentional characterization, but in others it will involve a challenge to a belief. Only the latter are of interest in the present discussion.

Let us return to the case where Willard and Wilfrid are discussing Willard's aesthetic preferences. In order for this discussion to move forward, many things must be shared taken-for-granteds. It must be taken for granted that El Greco was a painter, the meaning of "painter", what *Las Meninas* refers to, etc. But now it arises that Willard and Wilfrid do not share something. They do not both think that El Greco painted *Las Meninas*. This belief strikes Wilfrid as mistaken because of the context that has been built up to this point: the utterance "El Greco painted *Las Meninas*" only has significance relative to the forms of life that make these words mean anything for Wilfrid. Wilfrid thinks the belief problematic because it seems to violate a standard of the context in which he and Willard have embedded themselves. According to the rules of the forms of life that Wilfrid interprets Willard's belief to be following, the utterance is posited as a true statement that relates a painter to a painting. Only because Wilfrid understands

---

[24] Again, the evaluator would have to judge that the believer is acting reasonably similarly to herself, according to the explication of how this is done provided above.

the utterance as such does it seem unwarranted. The problem resides in a relatively narrow-scoped component of common ground, but a problem in any breadth of scope can bring a belief under scrutiny. For instance, stating something logically impossible is quite broadly unwarranted while stating something about the *Theaetetus* that every scholar has found reason to dispute is quite narrowly unwarranted. A broadly unwarranted belief will oftentimes be more difficult to interpret, because it is more difficult to accommodate something that seems obviously false. It may seem that the belief entered has been badly worded or even disingenuous. Thus, most of the epistemic action occurs in the relatively narrow-scope aspects of the various components that make up a context.

If the evaluator expresses her challenge to the believer, then the process of belief defense, or *justification*, starts. The internalist had argued that in acts of justification we see the true reasons that the believer held her belief. However, I think that what we really see is reasons that the believer believes the evaluator will accept. The process of belief defense, of justification, is not necessarily an unveiling of one's preexistent reasons for holding a belief—though that is a tactic that we employ at times—but rather a presentation of a series of potential reasons that a belief should be considered warranted, given what the believer takes the evaluator to agree with her on. The believer makes her epistemically responsible attempt to return to common ground. In deciding what the believer will agree with her on, the believer must rely on the same representations of conversational score and principles of humanity that I sketched above. The believer wants to be seen as epistemically responsible, so she presents reasons why her belief follows the rules of the justification game in which it was played. So, for instance, imagine that Wilfrid says that he actually thinks Velasquez painted *Las Meninas*. If Willard wants to defend the warrant of his original belief, he had better cite some source that he thinks Wilfrid will see as

a good indicator of the authorship of *Las Meninas*. What he thinks Wilfrid will count as a good indicator will depend on his interpretation of Wilfrid's representation of the score.

Justification will (ideally) continue until the evaluator becomes satisfied that the challenged belief is warranted. Foundationalists of various sorts have thought the only real stopping points were basic beliefs. If one truly named all one's reasons for holding a belief, one would eventually come to a foundation. My view is quite different, but will lead us into some familiar territory. To make it clear, I want to borrow a notion from Daniel Dennett. When talking about moral discourse, he appeals to the concept of "conversation stoppers", which are "a variety of uncertain and temporary equilibria...tending to accrete pearly layers of supporting dogma which themselves cannot [necessarily] withstand extended scrutiny but do actually serve on occasion, blessedly, to deflect and terminate consideration".[xcii] What he has in mind here can be generalized to epistemic justification. Indeed, justification stoppers, as I will call the modified concept, sound a whole lot like beliefs taken for granted on account of methodological necessity, or Wittgensteinian hinges.

This similarity should not be too surprising. Evaluators challenge believers because they think that believers have violated a rule that the two have taken for granted. Believers defend themselves because they believe that they have not violated any rules. In order for the believer to convince the evaluator that she has not violated any rules, the believer must appeal to things that she thinks will indicate to the evaluator that her belief follows the rules. A believer will effectively defend her belief if she appeals to things that they both take for granted, which are things that form the hinges of discourse. We should not, however, take this state of affairs as an indication that the believer unveils the pre-existent structure of beliefs when engaging in belief defense, as the internalist picture implies. We do, often, point to beliefs we already knew

supported the challenged belief, but some of our most useful insights (that is, new ways of thinking of things, not beliefs that already underlay our challenged belief) come in the heat of belief defense. When asked to defend a belief, we just as often rationalize as ratiocinate. It is possible to *discover* one's assumptions.

What makes justification stoppers "uncertain and temporary" is the fact that they vary across contexts. But, again, common ground is not fixed. Whether or not a given belief counts as a justification stopper depends on the participants; it stops justification if they share a justificatory point of view on it. This point of view could be that the challenge was answered satisfactorily, that the believer realizes she is unjustified, or that the participants realize that they just are not on the same page. A justification stopper might be terminatingly convincing or unconvincing.[25] The deflection and termination of justification is not "blessed", it is based on the profane fact that we humans need to see things similarly to engage in justification in the first place.

---

[25] Though I tend to think that the last option is merely a temporary state of affairs pending future opportunities for justification.

OBJECTIONS AND RESPONSES

Justification stops once the participants see eye-to-eye on what can be taken for granted or once they realize that they won't be able to do so. But couldn't the justification just keep going? Couldn't the evaluator just keep challenging potential reasons for a belief *ad nauseam*? Aren't we back in Agrippa's Trilemma? It seems that in principle two people with extraordinarily different points of view could find themselves unable to end justification, because they are unable to find common ground. Indeed, it seems that this shows up all the time in practice. Look at the disagreements between liberals and conservatives, for example.

Clearly, this possibility is unacceptable. Many think that only two ways of resolving it are possible. These responses are reminiscent of those of the two philosophical types that William James described as "tough-minded" and "tender-minded".[xciii] The first claims that some stopping points are inherently superior to others. Most often the superiority is justified by claiming that the stopping points that the tough-minded theorist has in mind more accurately represent reality. The tender-minded theorist, on the other hand, claims that no stopping point is better than any other. Indeed, she might go so far as to claim that no point of view on reality is better than any other, so we should happy adopt a pluralist multiculturalism where we go by the motto "believe and let believe".

The latter position seems to imply exactly the unfortunate conclusion that the initial objection found unacceptable. If no potential reason for belief is better than any other, then what is to stop justification between two people from very different epistemic backgrounds from continuing infinitely? Additionally, one could be guilty of any of the three seemingly faulty options in Agrippa's Trilemma and still be considered epistemically responsible according to this tender-minded theorist. Furthermore, it certainly *seems* that some reasons for a belief are better

than others. Trivial examples of bad reasons include: "because I said so", "because I believe everything my elders tell me", and "because the sun is blue". Supposing the tender-minded theorist could somehow dispatch with these, she would still need to deal with less trivial cases of explaining combustion in terms of phlogiston and explaining one's opposition to Obama's presidency in terms of his plot to destroy the US government from the inside.

The tough-minded theorist has a different set of problems. How can one explain the inherent superiority of some reasons over others in a non-question-begging manner? We always run the risk of simply pointing out the reasons that we value according to our present forms of life without explaining why this form of life has some inherent superiority. The most common attempt to avoid this problem is to say that some reasons (and some forms of life) better represent reality. This only pushes the problem back, though. For what makes some belief represent reality better than another belief? The "better" here implies a value judgment, which would require us to explain why we value this form of representation rather than another. Or, as Rorty, puts it, we can find "no independent test of accuracy of representation…no test distinct from the success which is supposedly explained by its accuracy", and without such a test, we cannot understand what an appeal to better representation amounts to.[xciv]

More careful critiques of these approaches have been undertaken elsewhere, but what concerns me in responding to the objection to which they respond is finding a middle way between the two. My proposal is this: justification stops through an appeal to reasons whose superiority is immanent to forms of life. When we encounter somebody who thinks about things quite differently from the way we do, the best thing to do is to find at what level we can find common ground and proceed from there. We certainly can find some set of standards in common, even if these are low-level components of common ground. Suppose that, for instance,

Liz says to Jean that she believes in ghosts. Jean does not believe in ghosts, so she challenges

Liz's belief. If either is to convince the other that her belief is warranted and the others mistaken,

then she must appeal to evidence/reasons for belief that the other will find epistemically

acceptable. That is, the evidence appealed to would have to be part of their common ground. Liz

and Jean will probably disagree on quite a few epistemic things: which other people's beliefs

about ghosts that can be trusted, for example. But ultimately there will be some agreed upon

standard against which Liz and Jean can test Liz's belief.

What Liz and Jean will likely have to appeal to is the *cash value* of Liz's belief; they will

have to "set it at work within the steam of [their] experience[s]".^{xcv} Another way to put this idea

is to say that Liz and Jean will have to evaluate the practical results of holding such a belief. By

"practical results" I do not mean whether the belief would make Liz better equipped to cope with

her mother's death or something like that; rather, I mean the implications the belief would have

in our epistemic interactions with the world. To determine the worth of a belief in ghosts, Liz

and Jean might examine what such a belief would suggest that we would encounter in certain

situations—we might see the image of a dead loved one or hear one communicating with us from

beyond the grave. Liz would have to specify which of these conditions she thinks show that

ghosts exist and Jean would have to specify which conditions she would find acceptable. If they

do not encounter such situations, or if these situations could be better explained (by their mutual

standards) by other beliefs, Jean should think that Liz has no good reason to believe in ghosts.[26]

What makes a given stopping point better than another is the extent to which the

evaluator and the believer find that stopping point convincing. A stopping point is effective only

---

[26] Obviously the situation here is much more complicated than these vague pronouncements can do justice to. There has been much ink spilled in the philosophy of science about what it means to test theories, and what I mean to point out here is merely that these sorts of issues are relevant not only for scientific interaction with the world, but also for everyday epistemic encounters therewith.

if it stops justification. This discourse-dependence makes stopping points horizontal, mutable, and open to challenge in different contexts. A reason that stops justification in one instance of belief defense may be the belief being justified in another. Pragmatism about warrant thus allows for virtuous circles in justification. "We can only make honest efforts to eliminate the sorts of defects we know our thinking risks, and perhaps to expand our conception of ways things might go wrong, so as to be on guard against other potential sources of error. The best we can achieve is always to some extent provisional and inconclusive, but that is no reason to succumb to the fantasy of an external validation".[xcvi] Eliminating defects involves challenging assumptions/stopping points. But to do so, we must maintain other assumptions.

The easy follow-up objection to my response so far is that I have merely pushed the problem back a step. Everything with my position is hunky-dory if believer and evaluator can find some common ground, but what if they have incompatible standards? To put it concretely, if Liz simply does not accept any of Jean's reasons for rejecting a belief in ghosts, it seems that pragmatism about warrant gives Jean no resources to respond to Liz's belief, and, as such, the pragmatist must count both Liz and Jean as epistemically responsible even though their beliefs are incompatible. Suppose Jean has even explicitly explained absurd and contradictory implications of Liz's belief and yet Liz still persists in holding it. Even then, the objector says, my pragmatism must allow Liz to maintain her belief in ghosts with a clear epistemic conscience.

Not surprisingly, I do not think that my pragmatism forces me to allow irrational believers to maintain their beliefs with clear epistemic consciences. Rather, I think that the price they pay for the ability to maintain their beliefs in the face of defeaters is *epistemic irresponsibility*, or irrationality. Such believers have not paid Putnam's aforementioned "price of

admission into the community of sanity".[xcvii] But doesn't this still let such believers off too easily? Aren't we just giving them their own little world of warrant without condemning their belief? Again, I think not. What we should say here is rather like what Wittgenstein said about somebody who doubts whether the earth existed 100 years ago: "I would not understand [him] for this reason: I would not know what such a person would still allow to be counted as evidence and what not".[xcviii] Pragmatism does not let an epistemically irresponsible believer get off scot-free to go live in her own world of warrant. It relegates the believer to a world where the things she says simply do not mean the same things that we mean when we say them. One cannot, for instance, claim to believe in ghosts and then reject all of our attempts to nail down a definition of ghosts without offering an alternative that matches up with our conceptions in some way. This would mean that she simply does not understand what we mean by ghosts, and hence cannot participate in any discussions thereof. The price of avoiding standards of discourse is the failure to mean anything by one's utterances.

Even worse than that, it will be hard to determine if what the believer purportedly believes even counts as a belief. After all, in order for us to successfully characterize a person's behaviors as a belief, those behaviors need to function similarly to beliefs with which we are familiar (to beliefs in forms of life that we know about). If a person rejects any attempt of ours to characterize her belief according to rules with which we are familiar, then we will have a hard time counting her purported belief as a belief. If this is not too high a price to pay for getting one's own world of warrant, I do not know what is.

There remains, however, the possibility that points of view widely divergent from today's common sense might actually provide new and better standards for warrant about a given topic. A person who seems to have absurd beliefs may actually turn out to have better ones that we do

for accomplishing our epistemic goals. Philosophers of science have explored this possibility extensively since Kuhn and Feyerabend. I do not wish to add to this literature here except to note that we should, as Feyerabend cautions, maintain openness to alternative ways of thinking, however absurd-seeming.[xcix] Doing so does not mean granting warrant to everything, but rather maintaining a sense of perspective. What seems absurd now could be mainstream in the future.

While this sort of complexity and uncertainty may be a bit uncomfortable, I would argue that it reflects the complicated and uncertain reality of the workings of belief. For better or for worse, all we have to go on in evaluating belief is what we currently think are the best things to believe and our reasons for those beliefs. These beliefs and these reasons could one day prove inadequate, but until we have reason to move past them we should follow our best standards, because they work best for us now.

The objector may think that I cannot get off this easily. Making standards for warrant depend on our activities leaves open a harrowing possibility. Imagine, for instance, a completely delusional community with a whole series of beliefs warranted by their standards. Would we have to count their beliefs as warranted even though delusional? Take the real cases of the ancient Greeks, for example: they were wrong about a whole lot of things that they thought they were right about. Should we say that they were warranted in their beliefs? Most pressingly, *we* could be this community. Without discourse-independent standards for warrant, we could be in a new sort of skeptical scenario where all of our warrant attributions are actually mistakes without our knowing it. Pragmatism leaves open the possibility of massive error on our part.

This worry is similar to worries about the aforementioned Rortian formulation that warrant is whatever our contemporaries let us get away with.[c] One way of objecting to this formulation is to say that many people's contemporaries let them get away with patently

unwarranted beliefs, which is the objection I put forth in the preceding paragraph. But one might also object slightly differently that many of Rorty's contemporaries did not let him get away with that very formulation, though *he* still thought it warranted. So clearly even the pragmatist accepts some standard independent of whatever people surrounding him allows. Either way, a discourse-independent standard is necessary.

Responses to this objection have come from one of two places. Some have posited the existence of external standards for warrant, while others have argued that warrant is just what we have been trained to think it is. Given my above cursory reasons for rejecting both responses, I want to follow John McDowell in attempting to navigate a middle way between this Scylla and Charybdis.[ci] This middle way has it that any discourse does have a set of standards of warranted belief independent of our knowledge thereof but that these standards are established by our language games and discourses.

As mentioned above, McDowell's argument relies on a generalization of Aristotle's notion of *phronesis*. In this generalization, when one gains the ability to use conceptual capacities, one becomes initiated into forms of life and their associated traditions, which have rules and standards that come along with them. To engage in certain forms of life—to use different types of concepts—is to play by a certain set of rules. If one does not play be the rules inherent in the forms of life that come with the concepts one is apparently using, then one simply does not count as using those concepts. In order to engage in any epistemic activity whatsoever, one must be engaged in some particular activity. Each activity brings with it a set of rules for how that activity works. Not following these rules merely makes it such that one does not participate in the activity (running the bases after whiffing at the baseball is some sort of activity,

but it certainly isn't baseball). Being epistemically responsible means following the rules

immanent in one's activity successfully.

In this way of thinking about things, though warrant *does* depend on our activities for its

existence and for its nature, there is a sense in which one can objectively be held to certain

standards for warrant. However, which standards one is being held to are relative (objectively

relative) to the form of life in which one participates. Forms of life are vaguely delimited and

horizonal, meaning that sometimes what standards one is being held to is not clear. Even so,

most cases are clear, and even unclear cases can be cleared up to some extent by using the

principle of humanity. How specifically this works is only suggested at above and would be

fodder for a more careful examination.

For now, let us stick to examining examples to clarify the point, if not in sufficient detail.

Suppose that rather than challenging Willard's belief that El Greco painted *Las Meninas*, Wilfrid

made the same error and counted Willard's belief as warranted. Let us imagine that Willard is

lecturing Wilfrid on art history, so Wilfrid has reason to trust Willard. We can say that Willard,

the believer, has made a mistake even though Wilfrid, the evaluator, has not caught it. We say so

not because some external standard exists that Willard has violated. Rather, Willard violates the

standard of the discourse of art history in which he is embedded—a man-made entity. Supposing

he actually believes that El Greco painted *Las Meninas* (and isn't merely guilty of a slip of the

tongue), he can be said to have been improperly initiated into the practice of art history. At

present Willard and Wilfrid operate with improper standards because, to awkwardly extent the

metaphor, they simply are not playing the game correctly.

In this way, discourses *do* have standards that do not depend on the specific participants

in the discourse, and thus meet the initial worry. We still have not addressed the skeptical

scenario, though. Suppose *we* were in the sort of situation in which Wilfrid and Willard found themselves with respect to standards of art history with some everyday belief of ours. This scenario does not worry me as much as it is supposed to. We *could* be operating with standards for warrant that ultimately prove unacceptable; in fact, this is likely the situation we are actually in. It may be the case that the goals we have for certain discourses are simply not being met by our present ways of engaging in those discourses. Still, we can always improve our standards of warrant. Though we now scoff at belief in its existence, phlogiston used to be the best means we had of explaining combustion. However, belief in phlogiston eventually died out because it did not explain everything that we wanted to explain with a theory of combustion embedded in a more general chemical theory. It did not explain all of the experimental results that we hoped our theory of combustion would explain. Our beliefs and our standards for belief evolve in concert with our interactions with the world.

Perhaps this situation is uncomfortable to some. I suggest that the problem does not fall out of a weakness of my theory but rather reflects an unfortunate fact about the way our discourses actually work. We just do not have perfect information about the world, the structure of our beliefs, or the relationship between our beliefs and the world. The best we can do is to continue to perfect our ways of thinking about the world by testing them against the lifeworld[27] and against other ways of thinking. We move forward in Neurath's boat because we are out at sea without any other alternative.

---

[27] "Lifeworld" is a Husserlian concept for the way a world presents itself to us when we engage in a particular form of life.

SKEPTICISM AND TRUTH REVISITED

Thus ends my preliminary positive picture of pragmatism about warrant. Because much

of it used terminology and concepts out of the mainstream of analytic epistemology, it may be

worth summarizing before moving on to some implications. The main proposal is this: we can

never take a view from nowhere in analyzing the warrant of a given belief of some believer. To

reflect on warrant attributions we cannot simply evaluate the believer's noetic architecture in the

neighborhood of the belief in question (whether this architecture be evidential, coherent, or

otherwise); we must also understand how the evaluator models the believer's noetic architecture,

even if this evaluator is us. Warrant evaluations are always an encounter of epistemic values. The

evaluator of warrant must interpret the reasons that the believer accepted a belief and then

evaluate her interpretation of these reasons.[28] These pragmatic readjustments to our

understanding of warrant analysis have definite applications in epistemological practice. I will

examine their implications for addressing external-world skepticism and for the role of truth in

warrant evaluations.

As mentioned above, internalism has long struggled with defending itself against

external-world skepticism (I will drop "external-world" henceforth). Work by Wittgenstein,

Williams, Williamson, and Chalmers (and, in different traditions but strikingly reminiscent,

Heidegger, Merleau-Ponty, and Deleuze)[29], among others, has provided reasons to think that its

defense woes are not merely incomplete arguments but rather betray deep problems with its

---

[28] By "reaons" I do not necessarily mean in the internalist sense. The reason somebody accepts a given belief could be, for instance, that they have been brainwashed.

[29] I will not get into these 20th-century continental reactions to skepticism here, because all of the set-up and exegesis necessary to do so would distract from the main points. However, for those interested, cf. op. cit. endnote xii for Heidegger; *The Visible and the Invisible*, chapter 1 section 1 for Merleau-Ponty; and *Difference and Repetition*, chapter 3 for Deleuze. Those familiar with these works should note that all three argue (as I do) that the traditional epistemological approach to skepticism, and not skepticism itself, is the problem.

approach. I will argue that the attitude adjustments engendered by pragmatism about warrant do much to alleviate skeptical worries.

We have seen that the Cartesian skeptic succeeds so well against the foundationalist merely by attacking the validity of foundational beliefs. But if one takes a step back from foundationalist orthodoxy, one realizes that the existence of foundational beliefs is a theoretical commitment peculiar to foundationalism. One must believe that a certain subset of beliefs supports every other belief in order to think that potentials defeaters for these beliefs are potential defeaters for an individual's entire set of beliefs.[cii] A related objection dooms coherentism. If the skeptic proposes that our beliefs, however coherent, may have nothing to do with reality, the coherentist must be worried specifically because her standards for warrant are internal and global. The worry is that even the coherentist's highest means of warrant granting would fail.[ciii]

Skepticism also attacks another weakness of internalist epistemology. If one believes that some possible fact about the world (e.g. that I am a brain in a vat) can elicit doubt about everything that she believes, then the skeptic has already won. Let us assume that the skeptic is right in saying that a skeptical belief can have nothing count *against* it, for that is how the skeptic has formulated it and that is what causes the anxiety in the traditional epistemologist. But what the internalist overlooks is that a skeptical belief also can have nothing count *for* it for the very same reasons. If we think that the only reasons we have for holding a belief are internal, then two beliefs that are internally equivalent are separated by a distinction without a difference.[30] There is, then, *no way to evaluate a skeptical claim*, and when the traditional epistemologist takes for granted that there is, she concedes to the skeptic in advance.

---

[30] Williamson argues that the skeptic (and the internalist in being threatened by the skeptic) assumes a phenomenalist version of evidence without argument. He then provides an alternative account of evidence that he thinks more plausible (Cf. chapter 8).

When we step back from internalism and its peculiar epistemological commitments, we see that in providing scenarios that are supposed to invoke skepticism, the skeptic has only succeeded in putting forth alternative metaphysical theories about the ultimate nature of reality. Perhaps, the skeptic proposes, reality is just a series of stimulations to my brain, which is in a vat somewhere. Fortunately, we do not normally take such theories to invalidate all of our beliefs at once. Metaphysicians propose new theories all the time without changing very many of their other beliefs. This is because we simply do not have the option of rejecting our entire set of beliefs in light of a new possibility (that, for instance, I am and always have been a brain in a vat), as a skeptical argument might suggest. It would be impossible to judge the warrant of the new possibility without any beliefs, for then we would be unable to determine what a reason to believe would be. We need a form of life with which to evaluate the new possibility. This epistemological idea is parallel to Quine's linguistic/ontological concept of a background language/theory. Our set of beliefs, our background theory, gives a new idea sense, because ideas can only have sense relative to our background theory. If there were no background theory by which to make sense of a novel belief, there would be no way of making sense of a novel belief. So we cannot drop all of our beliefs in favor of a new one.[civ]

But perhaps it seems that we do not need any special epistemological commitments to be taken in by skepticism.[cv] Anthony Rudd, for instance, points to the musings of a teenager untrained in philosophy to show that radical skepticism is far from a technical problem unique to professional philosophers.[cvi] But how one interprets such musings depends on one's epistemological commitments, even if the musings themselves may not. I would argue that such admittedly natural reflection on the nature of reality (to which I myself was prone as a teenager untrained in philosophy) is just that—ontological speculation—rather than doubt about the

warrant of all of one's beliefs. The interpretation of such reflection as a serious epistemological problem, not the reflection itself, is the uniquely philosophical problem. A further argument, and along with it some epistemological commitments, would be needed to turn innocent reflection into the worrying doubts of skepticism.

Hence, when presented with the possible belief that I could be a brain in a vat, I could accept it without rejecting *all* of my other beliefs. I would have to modify *some* of my beliefs about the underlying metaphysical nature of reality, but that would not affect the belief that, for instance, I had lunch at 12:30pm today.[31] Almost everything that we think about the world should still apply except for our belief about what makes the world that way. As William James has argued, "our knowledge grows in spots", not as a whole.[cvii] Currently, we turn to physics for our belief that the underlying structure of reality is stated by some version of quantum mechanics, but if some discovery arose to lead us to think that the underlying structure of the world is actually a very powerful mind or computer program[32] we would simply modify our beliefs about the underlying structure of the world. We would not, however, completely change the way we think about the world in every context. Belief in quantum mechanics has not made us change our belief that a table has legs, so why should the belief in a brain-in-a-vat ontology such as a computational structure of the world be any different?[cviii]

Whereas the internalist had problems with external-world skepticism, the externalist had problems basing normative principles on truth. Externalism went forth proposing qualities all warrant-producing faculties had in common without asking the prior question of why we think

---

[31] It might affect what it *means* to have lunch at 12:30, because we analyze having lunch as really just an epiphenomenon of…but that does not mean that it changes the fact that I did have lunch today at 12:30, as long as my original belief was correct.

[32] Both of these hypotheses have actually been proposed as metaphysical theories: the former by the German Idealists and the latter by Nick Bostrom, "Are You Living in a Computer Simulation?" *Philosophical Quarterly* 53 (2003) 243-255.

these faculties produce warranted beliefs. It seemed obvious that we should choose the faculties that produce *true beliefs*. This seemed a normative judgment hardly worth mentioning.[33] One problem, though, is that we are still in the process of figuring out what is true.[34] So, whereas externalism seems to work quite well for faculties that produce beliefs that we all agree are true (if only because we all agree), it fails to say much beyond these obvious cases. Perhaps more importantly, it fails to tell us *what it is about these beliefs* that make us think that they're all true. I have briefly argued that this failure merely leads the externalist to point to beliefs that the externalist values as true rather than using some principled criterion for selecting true beliefs. Plantinga and Goldman, for instance, value many overlapping beliefs, but certainly not all, as in the case of belief in God, and, not coincidentally, their different theories reflect their different values. Stephen Stich has argued more extensively that once we look at the details we find that all it is to call a belief true is to point at beliefs one has been acculturated to value.[cix] This will certainly not do.

Perhaps Ernest Sosa can be seen to have a better alternative. His virtue epistemology bears many similarities to Plantinga's proper functionalism, while acknowledging that warrant attributions are normatively governed behaviors. He proposes to evaluate belief-formation just like any other activity, by appealing to what we value in the exhibition of that activity. What we value is aptness. A skill is performed aptly if it succeeds because of the adroitness of the performer. In the doxastic case, we value a belief that is true (i.e. successful) and gained through

---

[33] Externalists are certainly not alone in using this technique. My objections are focused on my treatment of the externalist's picture, but they have relevance to all epistemologists who base evaluative epistemic judgments on truth.

[34] Plantinga takes the Cartesian tack of positing a benevolent God who allows us to have true belief. This is his guarantee that our faculties are in a congenial cognitive environment. An atheist or agnostic will find this methodology unacceptable, and I think that this means that we have no such guarantee. I agree with Stephen Stich (op. cit.) that we should not be too concerned that our beliefs are not guaranteed to be true, for both the reasons he offers and reasons I point to elsewhere in this essay.

some competence of the believer. That is, we value cases where a believer has discovered the truth of a belief through the proper means (means that we value as proper).[cx] One might think of Sosa as stating that we value beliefs acquired through a properly functioning faculty with the addendum that what counts as proper function depends on what we consider virtuous epistemic behavior.

This sort of virtue epistemology only goes part of the way, though. We are still left with a hand wave about truth. To the question of what we should believe and why, Sosa would presumably respond "apt beliefs, because they are true and gained virtuously". But what about cases where we don't know if our belief is true or if the very truth of the belief is in question? Worse, it seems that at least sometimes we *establish* the truth of a belief *by appealing to* its virtuous discovery. Sosa might say that in cases like this we can just appeal to adroitness, which is aptness minus truth/success. He notes that adroitness should be seen as having fundamental epistemic worth apart from and alongside truth.[cxi] But if this is the case, why point to truth at all? If we take Stich's and my arguments seriously, then to say we value true belief is to say that we value beliefs that we value. To say that we value apt belief is to say that we value beliefs that we value gained through means that we value. This analysis is not very helpful unless we understand how we determine which beliefs we value and why. That is exactly what I have tried to do in my positive account above.

If I am right, then we need not understand warrant as connected with truth and its difficulties. To understand the workings of warrant-attribution we need only understand how people go about evaluating each other's beliefs. This involves knowing what people's goals are for those beliefs, so an externalist might defend herself by saying that many people hope that their beliefs are true. To say that people want their beliefs to be true, though, is a vague

pronouncement. People may indeed value truth in evaluating each other's beliefs, but if we are to understand what a given person values when she seems to value truth we need to be very careful to take into account the differences across cultures and individuals concerning these values. Different people call different things true. This adjustment does not force us to believe, along with Dewey, that truth is just warranted assertion. Rather, it frees us to investigate warranted assertion without worrying about whether or not truth is at issue.

CONCLUSION

In order to highlight various implications of adopting a pragmatist account of warrant, I have generally favored breadth over depth. Some of this I supplemented by referring the reader to other work, but this is hardly satisfactory. After all, not all of the works referred to agree with each other—some even propound contradictory positions. Though drawing with such broad strokes does have such shortcomings, I think that the outline of a theory that I have sketched provides ample space for filling in and some guidance for how to go about it. Given the content of the pragmatist theory—the fact that it focuses on actual practices, cultural differences, language use, etc.—there are two general sorts of filling in that can be done. Armchair debate about the details of the conceptual framework would be helpful for parts of the picture. This approach would do well in specifying the role of occurrent beliefs, for instance, or in getting into more detail about the nature of truth and how different positions on its nature could affect the role it plays in justification and warrant attributions. I have hinted at my own positions on these issues and others, but much more philosophical work would need to be done to generate models that would be satisfactorily detailed. My outline also leaves room for filling in through empirical investigation. We could answer such questions as "how good are we at knowing what another person believes?", "how do we establish some statements/beliefs as acceptable?", "are there cultural differences in what beliefs a given perceptual experience generates?" through experimentation. All three of these questions have been addressed to some extent in the literature: psychological work on theory of mind relates to the first, psycholinguistic work on using language in order to accomplish an activity has given some insight into the second,[cxii] and work on perception in cultural psychology points to a "yes" answer to the third.[cxiii] Many more

such questions could be generated and answered by operationalizing aspects of the pragmatist theory of warrant I have proposed and putting them to the experimental test.

Recall once again that in inquiring into the nature of warrant, I have been using the question of what it takes for something to count as good to believe. The fact that, for a pragmatist, empirical results can bear on the answer to this question indicates something crucial about pragmatism about warrant. According to pragmatism (I will leave out "about warrant" for the remainder of the discussion), figuring out what it takes for a belief to be warranted is not a matter of mere conceptual analysis. We cannot simply assume that there is one correct parsing of warrant and/or justification; we must see how people actually go about justifying and judging the warrant of beliefs. In order to figure out what warrant means, pragmatism moves away from *analyzing the concept* of warrant and towards *investigating the use* of the concept. This might mean actually researching how people judge the worth of one belief versus another or how we can manipulate these judgments. It also means that we should not be alarmed when we find different people using different standards of warrant or the same person using different standards of warrant in different situations. This simply might be how belief evaluation works.

On the other hand, I think that the analyses of warrant that I spent the first part of this essay arguing against provide many useful insights on how we actually go about justifying and establishing the warrant of our beliefs. Like Earl Conee, I believe that the foundationalist, the coherentist, the externalist, and the internalist all have something to bring to the table in figuring out what is involved in our epistemic norms.[cxiv] Each theory has its own intuitive appeal, and for good reason. After all, what these philosophers have generally been considering in performing their conceptual analyses is their own usage of the concepts of warrant and justification. What they have done, then, is to mine in some detail their own norms for a range of cases. They have

provided *pro tanto* criteria for several of our epistemic practices. If we understand these criteria as part of a pluralist list of standards rather than as analyses of universal standards, then they can be quite useful. For instance, externalism has been and can still be important in understanding how immediate justification works (or, to put it another way, how basic beliefs gain their warrant). Much key work has been done by externalists to separate out basic beliefs from non-basic beliefs, to move the focus away from reasons towards norms, and to point out what sort of cognitive processes we value. Sosa's virtue epistemology has moved this process forward by highlighting more strongly these evaluation's connections to our normative practices.

There are many more important analyses to be found in these more conventional forms of epistemology. Pragmatism does not turn its nose at these issues, but rather reframes their relevance. If what I have had to say in this essay is correct, then we cannot merely take the analyses such as that of the externalist at face value. We have to ask, for instance, if what we currently count as properly functioning cognitive faculties should be considered properly functioning all things considered. And we should be open to the fact that our standards for proper function may change or evolve, because the way we conceptualize reality and our goals for a given discourse may change or evolve. We should also be open to others' standards and the possibility that they—not we—have better ones. The results of externalism are tools that the epistemologist can wield in examining certain aspects of belief evaluation. They can be used together with the tools provided by examining coherence relations and exploring the nature of evidence. We might even be able to classify different discourses according to how much they stress each of these epistemic values, for instance. Possibilities abound.

In all these details, my guiding question may have gotten lost, so let me ask it one more time. What is it for a belief to have warrant? To answer in short may not be helpful, but

hopefully the preceding will contextualize my response. For a belief to have warrant it must fit

the standards of the context in which it was offered. How these standards and contexts work is

complicated, but the foregoing has provided some preliminary thoughts on them. Most

importantly, I hope to have driven home the point that what makes a belief warranted depends on

a relationship between the warrant evaluator and the believer rather than on objective facts about

the believer. Does my view help us decide which beliefs to accept? Perhaps. My vague

recommendation is that we accept those beliefs that seem the best based on our current standards.

Though this does not say much, it does imply that we can be warranted now even though

ultimately mistaken and that we should not assume that our current standards for warrant (no

matter how precise or well-worked-out) are the only or the best. However, a more detailed

prescriptive picture will have to be saved for another discussion.

---

[i] Plato, *Theaetetus*, trans. Benjamin Jowett (Internet Classics Archive), http://classics.mit.edu//Plato/theatu.html (accessed September 29, 2009)

[ii] Edmund Gettier, "Is Justified True Belief Knowledge?" *Analysis* 23 (1963): 121-123.

[iii] Timothy Williamson, *Knowledge and Its Limits* (New York: Oxford, 2000)
See also David Lewis's "Elusive Knowledge" *Australasian Journal of Philosophy* 74 (1996)

[iv] Michael Williams, *Unnatural Doubts* (Princeton: Princeton University Press, 1996) passim, esp. Ch. 3
See also his "Skepticism" in *The Blackwell Guide to Epistemology* ed. John Greco and Ernest Sosa (Massachusetts, Blackwell, 1999)

[v] See Michael Williams, *Groundless Belief* (Princeton: Princeton University Press, 1999/1977) for a thoroughgoing rejection of many of these reasons.
See also Roderick Chisholm "A Version of Internalist Foundationalism" in *The Foundations of Knowing* (Minneapolis: University of Minnesota, 1982) for reworked internalist reasons.

[vi] Cf. James Pryor, "There is Immediate Justification" in *Contemporary Debates in Epistemology* ed. Matthias Steup and Ernest Sosa (Massachusetts: Blackwell, 2005) for such an ambiguous position.

[vii] W.V.O. Quine, "Two Dogmas of Empiricism" *The Philosophical Review* 60 (1951): 20-43

[viii] E.g. Laurence BonJour, *The Structure of Empirical Knowledge* (Cambridge: Harvard University Press, 1985), Ch. 5

[ix] Op. cit.

[x] On concepts as metaphors, see Nietzsche, "On Truth and Lying in an Extra-Moral Sense" in *Nachlass* trans. Walter Kauffman and Daniel Breazeale
On narratives as explanatory, see James Brown, "Explaining the Success of Science" *Ratio* 27 (1985): 49-66
On metaphors as guiding our intellectual life, see George Lakoff and Mark Johnson, *Metaphors We Live By* (Chicago: University of Chicago, 1980)

[xi] Williams, *Unnatural Doubts*, 109.

[xii] Immanuel Kant, *Critique of Pure Reason* trans. Paul Guyer and Allen W. Wood (Cambridge: Cambridge University Press, 1998): Bxxxix. Cited in Martin Heidegger, *Being and Time* trans. Joan Stambaugh (Albany: State University of New York Press, 1996): 189 in the midst of an informative discussion of skepticism.

[xiii] Williams discusses and refutes these attempts in all three works of his that I have already cited.

xiv Most of the explication of reliabilism comes from Alvin Goldman, "What is Justified Belief?" *Epistemology: The Big Questions* ed. Linda Martín Alcoff (Massachusetts: Blackwell, 1998): 89-109.

xv Alvin Plantinga, *Warrant and Proper Function* (New York: Oxford University, 1993): viii-ix

xvi Stephen Stich, *The Fragmentation of Reason* (Cambridge: MIT, 1990): Ch. 4.6.2

xvii ibid, 92

xviii In *Warrant and Proper Function*, Ch. 12

xix For a discussion of jettisoning the notion of truth in favor in instrumental value, see ibid, Ch. 6

xx Op. cit. and Alvin Plantinga, *Warranted Christian Belief* (New York: Oxford University, 2000): 178

xxi *Warrant and Proper Function*, 13

xxii Wilfrid Sellars, *Empiricism and the Philosophy of Mind* (Cambridge: Harvard University, 1997): 74, originally published in *Minnesota Studies in the Philosophy of Science* 1 ed. Herbert Feigl and Michael Scriven (Minneapolis: Minnesota University, 1956)

xxiii ibid, 75

xxiv ibid, 79

xxv For a related argument, cf. Paul Feyerabend, *Against Method* (New York: Verso, 2008/1975): Ch. 9

xxvi Cf. Earl Conee and Richard Feldman, "Evidentialism" in *Evidentialism: Essays in Epistemology* (New York: Oxford University, 2004): 83

xxvii Earl Conee "First Things First" in *Evidentialism: Essays in Epistemology*, p. 35

xxviii Conee and Feldman "Internalism Defended", same volume, p.79.

xxix Richard Feldman, "Having Evidence", same volume, 226.

xxx Cf. "Evidentialism", 85

xxxi ibid, 94

xxxii See Stephen Stich "Normativity and Epistemic Intuitions" in *Philosophical Topics* 29 (2001): 429-460 for some of his experimental results and citations of others'.

xxxiii Cf. Stich, *The Fragmentation of Reason*, Ch. 4 (esp. section 4.6) for detailed arguments for the arbitrariness of our epistemic evaluations. Also see the first three chapters of the same volume for a series of arguments for the irrationality of our cognitive faculties.

xxxiv Ludwig Wittgenstein, *Philosophical Investigations* trans. G.E.M. Anscombe (London: Prentice Hall, 1958): #325

xxxv Hans-Georg Gadamer, *Truth and Method* trans. William Glen-Doepel, revised by Joel Weisheimer and Donald G. Marshall (New York: Continuum, 2004/1975): 295, italics added

xxxvi Ibid, 291

xxxvii Cf. ibid, 296, 331

xxxviii Donald Davidson, "On the Very Idea of a Conceptual Scheme" *Proceedings and Addresses of the American Philosophical Association* 47 (1973-1974): 5-20, passim, esp. 18-19.
Also see Donald Davidson, "Radical Interpretation Interpreted" *Philosophical Perspectives* 8 (1994)

xxxix Herbert H. Clark, *Using Language* (Cambridge: Cambridge University, 1996): 92

xl Ibid. Citing Steven Levinson "Activity Types and Languages"

xli Ibid, 62

xlii Cf. ibid, 94 and David Lewis, *Convention: A Philosophical Study* (Cambridge: Harvard University, 1969): 53

xliii *Philosophical Investigations*, #18

xliv Ibid, #23, his italics

xlv Ludwig Wittgenstein, *On Certainty* ed. G.E.M. Anscombe and G.H. von Wright, trans. Denis Paul and G.E.M. Anscombe (New York: Harper Torchbooks, 1969): #358

xlvi Ibid, #110

xlvii For his development of this concept, see *Being and Time*, Sec. 13

xlviii Edmund Husserl, *The Crisis of the European Sciences and Trascendental Phenomenology: An Introduction to Phenomenological Philosophy* trans. David Carr (Evanston: Northwestern University, 1970): 111

xlix "Accomplishments can only be understood in terms of the activity that accomplishes them" (ibid, 117).

l Williams, *Unnatural Doubts*, 117-18

li Sellars*Empiricism and the Philosophy of Mind*, 79

lii Cf. *On Certainty* #105

liii *Philosophical Investigations*, #241

liv *On Certainty*, #103

lv Triandis "The psychological measurement of cultural syndromes". *American Psychologist*, 51 (1996): 407-415.

lvi Markus, Kitayama, and Heiman "Culture and 'Basic' Psychological Principles". *Social psychology: Handbook of basic principles* (New York Guilford Press, 1996): 857-913

lvii On the importance of sharing of common ground, see Clark, op. cit., 94-100. He also provides some categories for things that need to be shared later in this chapter.

lviii John McDowell, *Mind and World* (Cambridge: Harvard University: 1993): 84

lix ibid, 82

lx *Philosophical Investigations*, #198

lxi Richard Rorty, *Philosophy and the Mirror of Nature* (Princeton: Princeton, 1979)

lxii Cf. Stich, op. cit. 140-149

lxiii Adam See, "Wittgenstein on the Essence of Grammar." Weblog entry. The Space of Reasons. May 2, 2009. May 2009. (http://thespaceofreasons.blogspot.com/2009/05/wittgenstein-on-essence-of-grammar-adam.html).

lxiv *Philosophical Investigations*, #373

lxv Op. cit., 91-5, 109-110, et passim

lxvi Hilary Putnam, "Why Is a Philosopher?" in *Realism With a Human Face* ed. James Conant (Cambridge: Harvard, 1990): 114

lxvii *Truth and Method*, 318

lxviii *Philosophical Investigations*, 202

lxix Crispin Wright, "Rule Following Without Reasons" *Ratio* 10 (2007): 15

lxx Ferdinand de Saussure, *Course in General Linguistics* ed. Charles Bally and Albert Sechehaye, trans. Wade Baskin (New York: McGraw Books, 1959).
See also W.V.O. Quine, "Two Dogmas…" op. cit.

lxxi Feyerabend, op. cit., 272

lxxii *Philosophical Investigations*, #83.

lxxiii *The Crisis of the European Sciences*, 149

lxxiv Op. cit.

lxxv Jürgen Habermas, *The Theory of Communicative Action, Vol. 2: Lifeworld and System: A Critique of Functionalist Reason* trans. Thomas McCarthy (Boston: Beacon Press, 1987): 122 and 123

lxxvi Habermas, op. cit., 124

lxxvii David Lewis, "Scorekeeping in a Language Game" *Journal of Philosophical Logic* 8 (1979) 339-359.

lxxviii Ibid, 342-344

lxxix Ibid, 349

lxxx Ibid, 345

lxxxi Ibid

lxxxii Clark, op. cit., 103-116

lxxxiii Op. cit.

lxxxiv Op. cit., 346

lxxxv Cf. Clark, 104

lxxxvi Lewis, op. cit., 347

lxxxvii Davidson, "Conceptual Scheme", 19

lxxxviii Ibid, 18

lxxxix Grandy, "Reference, Meaning, and Belief", 443, cited in Stich, 46

xc Stich, ibid, 47

xci Cf. Williams, "Skepticism", 50 et passim

xcii Daniel Dennett, *Darwin's Dangerous Idea* <finish citation>: 507

xciii William James, *Pragmatism* (Toronto: Dover, 1995/1907): 12

xciv Richard Rorty, "Introduction: Antirepresentationalism, Ethnocentrism, and Liberalism", *Objectivity, Relativism, and Truth* (Cambridge: Cambridge University, 1990): 5

xcv James, op. cit.

xcvi McDowell, op. cit.

xcvii Op. cit.

xcviii *On Certainty* #231

xcix *Against Method* <cite pages>. "How to Defend a Society Against Science", available as an open-source online.

c Op. cit..

[ci] Cf. op. cit. passim

[cii] Williams, *Unnatural Doubts*, passim, esp. Ch. 3 on epistemological realism.

[ciii] For a much more extensive treatment of coherentism with which I am largely sympathetic, see Chapter 7 of Williams, ibid.

[civ] Quine, "Ontological Relativity" in *Ontological Relativity and Other Essays* (New York: Columbia University, 1969) 48-51

[cv] Barry Stroud and Anthony Rudd have both suggested this possibility.
Cf. Stroud, Barry, "Epistemological Reflection and Knowledge of the External World". *Philosophy and Phenomenological Research* 56 (1996): 345-358.
Rudd, Anthony. "Natural Doubts". *Metaphilosophy* 39 (2008): 305-324.

[cvi] Rudd, ibid, 312-313

[cvii] James, op. cit., 64. The discussion immediately following this quote is a highly lucid way of thinking about what I am arguing for here.

[cviii] cf. Chalmers, "The Matrix as Metaphysics" in *Philosophers Explore the Matrix* ed. C. Grau (New York: Oxford University, 2005), passim, especially 10-13

[cix] Op. cit.

[cx] Sosa, *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume 1* (New York: Oxford University, 2009)

[cxi] ibid, 88

[cxii] See Clark, op. cit. for discussion and citations

[cxiii] See, for instance, R.E. Nisbett and Y. Miyamoto "The influence of culture: Holistic vs. analytic perception" *Trends in Cognitive Sciences* 9 (2005): 467-473.

[cxiv] "The Basic Structure of Justification", same volume as previously cited